

BRITISH TELECOMMUNICATIONS ENGINEERING

Included in this Issue

Start of New Series:

Collaborations

Local Area Mobility

Fighting the Information Terrorists



**The Journal of The Institution of
British Telecommunications Engineers**



BRITISH TELECOMMUNICATIONS ENGINEERING

Contents

VOL 16 ■ PART 4 ■ JANUARY 1998

Guest Editorial: New Challenges, New Opportunities	265
Brian Spiegelhalter	

Collaborations

Switzerland—Life in a New BT Joint Venture	266
R. S. Murray Trantor	

Internet and Multimedia

BT WebWorld™—Web Hosting for Every Business Need	273
Ian Cameron and David Patrick	
Games People Play—The Wireplay Experience	279
Kelvin Lawrance, Dave Cleland, Bryan Johnson and Kevin Welsby	
A High-Performance Network for Gaming—Wireplay	284
William Millar, Phil Ashby, Martin Duffy and Kevin Welsby	

Mobility and Radio Communications

Local Area Mobility	290
Steve Buttery, Bob Merrett, Adrian Pote and Andy Sago	

External Engineering

The Underground and Copper Access Network Infrastructure	296
Mazz Da Rios	

Telecommunications in the 21st Century

Fighting the Information Terrorists	300
Simon J. D. Phoenix and John Regnault	
Beyond the Handset: Audio for the Next Millennium	306
Andrew Rimell and Mike Hollier	
A Wearable Communication Space	311
Mark Billinghamurst, Jerry Bowskill and Jason Morphet	

Interacting Demands for Multimedia Services	318
Ann Matthews	

Creating the Corporate Soloist	324
Who plays the Stradivarius?	

Mick Cope	
Closure of Cardiff Empire Exchange	337
Michael Clemitson	

Regular Features

pcochrane@btlabs	344
Book Reviews	345
Telecom Focus	348
Information	352

Cover Picture

Engineers working on a telecommunications tower based in Norway.
BT Corporate Picture Library;
a BT photograph

Board of Editors

Chairman Bill Mills, M.A., C.ENG., M.I.E.E.

Managing Editor Paul Nichols, B.SC.

Secretary Les Brand, M.ENG., M.B.A., D.I.C.

Treasurer Denis Taylor, F.C.M.A.

Assistant Editor Richard Kemp, B.SC., C.ENG., M.I.E.E. Roger Blake, C.ENG., M.I.MECH.E. Jason Chilton Martin Evans, B.SC., PH.D.

David Greenop, B.SC., M.SC. Peter Howell, B.SC., C.ENG., M.I.E.E., M.C.I.B.S.E., M.I.O.M. Peter McKenzie, B.SC., C.ENG., M.I.E.E., M.B.I.M.

Graham Neilson, B.SC. Murray Trantor, B.SC., C.ENG., F.I.E.E., D.M.S., M.I.MGT. Keith Ward, C.ENG., F.I.E.E., M.I.MGT.

David White, M.I.P.D., M.B.I.M., L.I.A.A.


Theme Editors

Collaborations Brian Spiegelhalter **Internet and Multimedia** Bob Foster **Mobility and Radio Communications** Nigel Wall

External Engineering David Crampsey **Telecommunications in the 21st Century** Peter Cochrane and Steve Sim

Brian Spiegelhalter

New Challenges, New Opportunities



It gives me great pleasure to introduce a series of articles from the joint ventures (JVs) that BT is establishing in Europe and beyond. The strategy to attack the £141 billion (by 2001) European telecommunications market has brought a new set of engineering challenges and consequent rewards. The articles, the first of which begins on p. 266, will look at establishing new JV telecommunication businesses which will initially challenge the incumbent PTT. To put this into context, I will briefly outline why BT is in Europe and how we are approaching the opportunity.

Continental Europe can be regarded as home territory for BT and has been a significant element of BT global strategy since the formation of Concert in 1993. However, with the growing European liberalisation, BT has a unique opportunity to attack the domestic fixed and mobile markets in the countries with the highest value telecommunications revenues. The approach we have taken is to form JVs with major indigenous companies so that local, rather than British, companies are created.

BT has JVs in Sweden (Telenordia), Italy (Albacom), Germany (Viag Interkom), Netherlands (Telfort), France (Cegetel), Switzerland (Newtelco) and Ireland, and has alliances with Norway PTT (Telenor), Denmark PTT (Teledanmark), Austria PTT, Portugal PTT (PT) as well as subsidiaries in Belgium and Spain and distributors in many of the Eastern European countries.

As the countries have liberalised and licences have been granted, the JVs have broadened their portfolios, initially through the distribution of the Concert global product range and lately with the addition of domestic value-added products such as Frame Relay Service and Virtual Private Voice Networks.

In most of the EU countries, 1 January 1998 marked the full liberalisation of public voice services (PSTN). The JVs in our European group are now launching services to take advantage of this opportunity.

To challenge the incumbent PTTs successfully, we are also seeking to differentiate ourselves by entering the rapidly evolving Internet market and acquiring mobile licences. We are not just setting up mobile services but, more importantly, converging the fixed and mobile business to provide a new range of integrated fixed and mobile services.

As you can see, we have been extremely busy in Europe over the last three years and we have encountered several challenges along the way. These include regulation, local ways of doing business, tight financial targets and differentiating ourselves from the PTTs and other competition.

Some of the engineering challenges are like those encountered in BT when introducing new technologies but others are very different. We have started from a green-field situation and are designing, planning and implementing networks and systems simultaneously to meet demanding service launch dates. In addition, we are developing innovative local access techniques as we do not have any legacy infrastructure. All this must be done in a new multicultural environment where some key skills are in short supply; for example, Internet and mobile.

BT is normally the main technical partner in the JV and will wish to place its own people in the key technical positions. In addition, we are expected to provide technical leadership by supplying expertise and experience from BT, driving innovative use of technology to differentiate from

the competition and driving consistent cost-effective technology across the set of JVs. This last point is commercially rather than architecturally motivated as consistent technology helps reduce costs (for example, better purchasing power, utilising shared developments to ensure seamless interworking between JV platforms and those of other JVs and Concert) and it enables rapid introduction of services developed in BT and other JVs.

However, BT's primary aim is that each venture is commercially successful and so we tread a careful path between exerting influence, but not interfering with the day-to-day running of the business.

Our engineering role falls into two parts: as a shareholder driving the performance of the JV and attempting to resolve any issues which may impede this, and as a supplier providing networks and systems resources, technical expertise and complete solutions to the JV and recovering costs through agreed commercial frameworks. The two roles are complementary and form the basis of our engineering relationship with the JVs.

Europe is more than a set of JVs: the whole is more than the sum of the parts. We are realising significant synergies by bringing the JVs together to share experience, developments and products. In addition there are opportunities for pan-European products and services.

Europe has opened up opportunities not only for BT's business but also for BT's engineers. We now have over 200 BT engineers working in or supporting the JVs in Europe. The engineering work has been challenging but stimulating and it has demanded close teamwork within the JV and between JVs, suppliers and partners. We have established a sound foundation to BT's vision of being No. 1 in Europe and are strongly positioned to maintain our engineering lead into the next millennium.

Brian Spiegelhalter
BT's Chief Engineer Europe

Switzerland—Life in a New BT Joint Venture

BT's global strategy continues with another joint venture in Switzerland. This article describes the hectic first few months of getting the company and its network up and running. It also outlines the opportunity and challenges for BT and its people.

Introduction

This article describes the challenges and opportunities for both BT and its people in the field of joint ventures (JVs) in what is a rapidly changing environment created by the liberalisation and globalisation of the telecommunications industry. Although every JV is different in many ways, there are common elements and benefits to be gained by collaboration. The article is based on the first 10 months of Newtelco, a BT joint venture in Switzerland from the perspective of the Chief Operations Officer, responsible for the network and company IT systems.

Characteristics of Switzerland

Switzerland has a total area of 41 284 square kilometres of which 26 651 square kilometres is productive. This is divided 41% forests, 35% meadows and arable land, 21% pastures and 3% orchards and vineyards.

As well as the well-known traditional industry of watchmaking,


Switzerland has big chemical and pharmaceutical, machinery manufacture and confectionery sectors. Although not a member of the European Community it does trade heavily with its members with approximately one quarter of its exports and imports being with Germany alone.

The resident population in 1994 was just over 7 million including 1·35 million foreigners. Major population centres are Zurich (nearly 1 million), Geneva (437 000), Basel (407 000), Berne (331 000), Lausanne (292 000), Lucerne (181 000), St Gallen (136 000), Winterthur (116 000) and Lugano (109 000).

Four main languages are spoken in Switzerland: German (64%), French (19%), Italian (8%), Romansh (1%) and other 7%. The distribution is shown in Figure 1.

There is a federal system of Government with each of the 28 cantons being autonomous. Within each canton there are communes or municipalities, each of which has its own income tax rate. Community decisions are often taken by referendum of the whole population.

Figure 1—Switzerland—languages and main cities



Newtelco AG was inaugurated on 15 April 1997 with the purpose of becoming a full second telecommunications operator in competition with the Swiss Telecom PTT (now renamed Swisscom).

The Birth of the Company

The legal name of the joint venture company is Newtelco AG and it was constituted on 15 April 1997. Since September 1997 products have been sold under the brand name of sunrise netcom services and this is growing in parlance as synonymous with the company name (see Figure 2 for the new logo).

Newtelco AG was formed by an agreement between:

- Schweizerische Bundesbahnen (SBB), the state-owned railway;
- Schweizerische Bankgesellschaft (SBG but now known as *UBS*), a big Swiss bank;
- Migros-Genossenschafts-Bund (MGB), the largest Swiss retail outlet;
- Tele Danmark; and
- BT.

The three Swiss shareholders own the majority of the company and it

Figure 2—sunrise logo



had an initial capitalisation of 210 million Swiss francs (£91 million).

The Shareholders' Agreement

The Shareholders' Agreement is a document some 120 mm thick and is written in English and German versions, with Swiss law applying. It records the different contributions from each shareholder and ongoing cooperation among them.

The company was inaugurated officially on 15 April 1997 with the purpose of becoming a full second telecommunications operator in competition with the Swiss Telecom PTT (now renamed *Swisscom*).

The agreement records that Newtelco will distribute Concert products as long as they meet commercial and quality-of-service criteria, and that BT and Tele Danmark (TD) will provide chargeable technical and marketing consultancy. The three Swiss partners agreed to consider outsourcing their internal networks to Newtelco, and Migros agreed to act as a retail outlet for the company's products.

SBB agreed to give Newtelco a wayleave to run cable alongside the railway tracks in return for shares. It also agreed to lay the cable, having been purchased on Newtelco's behalf, at reasonable cost and to maintain it thereafter. Newtelco can also lease dark fibre and transmission capacity from the internal SBB network and vice-versa. SBB also agreed to lease

some of its premises for Newtelco to install telecommunications equipment.

The BT Switzerland Division along with access to the BT/Concert portfolio of international products was contributed to Newtelco by BT in return for shares.

Start-up money for the new company was contributed by TD.

SBB and UBS contributed WANDA (Wide Area Network Driven by ATM), a previous network services joint venture between the two companies involving third-party customers run under a special licence from the regulator, BAKOM.

Moving from Company Formation to Physical Implementation

Prior to the launch of the company in mid-April 1997, a steering committee consisting of senior managers from each partner was formed to organise the legal launch of the company and to start planning the network implementation.

The steering committee sponsored the production of a network master plan which recommended the roll-out of a core fibre network over three years. One of its first tasks was to plan and order the fibre cable so that Newtelco could be first to market. The committee also appointed the senior management team of Newtelco as shown in Figure 3. The first stage of the issue of an invitation to tender for network equipment was completed in mid-March 1997.

Figure 3—Newtelco senior management




Figure 4—Newtelco backbone optical-fibre cable roll-out, 1997

A new board was elected with two members each from SBB,TD and BT, Steve Andrews and Geoff Almeida, and one from Migros and a Swiss lawyer; and this body now meets every six weeks with an Executive Committee, a sub-set of the Board, meeting midway between Board meetings.

From March onwards a lot of work had to be done in a very short time. Long-term employees had to be recruited on a large scale, vendor responses had to be evaluated, new equipment and office premises had to be searched for and a robust project plan to enable products to be launched on schedule had to be developed.


Office automation and IT systems had to be scoped and evaluated, the first of which was the billing system.

The Market and Competition

Although Switzerland is not a member of the EU it tends to follow EC legislation fairly closely and has agreed to open its state-owned telecommunications monopoly to competition on all services on 1 January 1998. At present, regulation is carried out by BAKOM but a Communications Commission (KomKom) has recently been appointed which is independent from government but has no legal powers until 1 January 1998.

Swisscom has a turnover of 10 500 million Swiss francs (£4600 million), has a workforce of 21 200 employees operating 4.6 million telephones, 58 000 payphones and 1 million mobile telephones. It also collects the radio and television broadcast fees and owns a third of Cablecom, the biggest cable TV company in Switzerland with a market share of over 50%. It is presently organised in 17 geographical regions but is in the process of restructuring. At 64 lines per 100 inhabitants the fixed network is fairly mature and the last of the analogue exchanges should be modernised by January 1998. From 1 January 1998 it will be 49% privately owned.

With a high gross domestic product and a lot of international trade



Switzerland is a good market for a new telecommunications operator to enter:

Apart from Swisscom, competition against Newtelco comes from:

- DIAX, a consortium of six electricity-generating companies and South Western Bell Communications (SBC) which has strung optical-fibre cables along the electricity pylon network;
- Global One, the international JV of Deutsche Telekom, France Telecom and Sprint from the USA; and
- Worldcom, a USA company strong on data.

As a complete second operator Newtelco is well placed to compete in all market segments but will naturally target those with high margins first. A turnover of 1000 million Swiss francs (£435 million) is forecast for the medium-to-long term with a market share of around 15%.

The Physical Network

Optical-fibre cables

As mentioned above, one of the first tasks was to design, plan and lay the optical-fibre cables along the railway.

This was originally planned as a 20-fibre network of roughly 700 km between the major cities of Zurich, Geneva, Basel, Bern, Fribourg, Luzern, Lausanne and Biel, as shown in Figure 4. For this year, Lugano, in the Italian sector, will be reached by the use of fibres leased from SBB. The figure-of-eight design gives some resilience and this is planned to be further improved when more cities are brought onto the network in years two and three when another 500 km each year will be laid. This will allow the present single-tier network to be configured into a three-tier ring-based structure as the network grows. The cable is all singlemode fibre of two different grades to allow upgrade of the bandwidth carried as the terminal electronic technology is developed. Eventually single-fibre working and wave-division multiplexing (WDM) should be possible.

One of the early modifications to this master plan was to lay some sections of the cable near to city centres with a higher number of fibres in order to prepare for some segments of local access rings in the future, as shown in Figure 5. This should lessen the future need to take up space in the SBB concrete channels, which is at a premium.

Figure 5—Concept of tapered fibre roll-out




Figure 6—Main node site in Zurich

At the time of writing, 95 per cent of the cable had been laid and splicing was in progress. The schedule was to have the complete network handed over and tested by mid-December in time for the 1 January 1998 launch. This is some three months behind the original plan and some expedients have had to be taken which will be outlined in a later section.

Accommodation

One of the biggest headaches faced has been the acquisition and fitting out of accommodation. In the absence of marketing data in the granularity required, searches were concentrated on SBB sites which were as near the city centres as possible. In some of the smaller locations existing SBB transmission rooms have been used, at least as an interim measure. The biggest sites are at Zurich and Lausanne and these were refurbished by Siemens. This solution was chosen as there is an extremely lengthy planning permission, detailed design and reconstruction process in Switzerland. This can take six months from start to finish but Siemens was in a position to complete this work in record time and to the very tight project plan. Other sites, after having been identified by Newtelco people, were let to a managing agent, who undertook detailed designs according to Newtelco specifications and employed local building contractors to carry out the work. (Figure 6 shows the main node site in Zurich.)

The original lease agreement with SBB did not fully meet the needs of Newtelco so this had to be renegotiated, which again injected unexpected delays.

Nevertheless, the process seems to have worked fairly well but a lesson for future projects is to start the accommodation process very early in the plan.

Voice platform

An invitation to tender was issued to eight vendors in mid-March 1997 based on the network master plan



that had been written by the original project team. This asked for replies by mid-April and five replies were received. Senior management then short-listed the manufacturers in early May and further intensive commercial and technical negotiations continued with Alcatel, Siemens, Ericsson and Nortel. All had different packages of price and services offered, and extensive evaluation based on a number of criteria had to be done. As a result, in early June the board accepted a recommendation from management to enter a sole supplier agreement with Siemens to cover the core network voice, transmission and data requirements.

The voice network consists of two Elektronisches Waehlsystem Digital (EWSD) switches which carry the service switching point (SSP) functionality. One service control point (SCP) based on Siemens Nixdorf hardware was provided to deliver intelligent network (IN) services. At

present the collection and aggregation of voice traffic is either by direct connection at 2 Mbit/s for primary-rate interface PBX connections, via 2 Mbit/s add-drop multiplexer (ADM) ports from the synchronous digital hierarchy (SDH) network or by indirect access via the Swisscom interconnect route (see Figures 7 and 8 for examples of the equipment).

Both switches are managed from the network management centre (NMC) in Rumlang, Zurich by the Node Commander proprietary system from Siemens.

Transmission platform

The transmission platform, also from Siemens, gives SDH functionality at STM-1, STM-4 and STM-16 levels. At present, there are no digital cross-connects planned, grooming being done within the add-drop, cross-connect and terminal multiplexers.

The transmission system is again managed centrally from the NMC

Figure 7—Siemens EWSD



Figure 8—Siemens SCP (foreground) and SDH (background)

using the Equipment Management Operations System (EMOS) proprietary management system.

Data platform

The data platform is again supplied by Siemens through a joint agreement with Newbridge, whose MainStreet family of products is being used to provide end-to-end service management for Internet, frame relay, asynchronous transfer mode (ATM) and time-division multiplex (TDM) customers.

The MainStreetXpress 46020 system is used for network management from the NMC.

Because of the delays in getting the SBB cables in and working, and refurbishment of some of the accommodation, the data network had to be launched in five phases:

- *Phase 1*—Use of the existing WANDA network from the two partners UBS and SBB. This is based on a network of four Cascade nodes interlinked by 34 Mbit/s links leased from SBB. This carries frame relay, local area network (LAN) interconnect and Internet service providers' traffic.
- *Phase 2*—Use of a leased line to connect Newbridge nodes at Zurich and Geneva.
- *Phase 3*—Use of SBB leased lines and Newtelco transmission equipment to connect seven nodes.



- *Phase 4*—Use of Newtelco fibres and transmission equipment to connect 20 points of presence (PoPs) within Switzerland.
- *Phase 5*—Migration of existing customers from the WANDA network to the Newtelco network.

The above programme enabled the data products to be launched according to plan.

IT systems

The requirements for IT systems and office automation for a small telecommunications operator the size of Newtelco are very different from either parent—BT and TD. Finding appropriate solutions is a challenge in itself. Many software houses have what they think are 'tailored' and 'flexible' solutions but all have down sides.

During the spring and early summer of 1997 an extensive internal customer-requirements capture exercise was undertaken. This was then used as a basis for system selection.

The installation of Oracle Financials, which had started under the auspices of BT Switzerland, was completed.

The top priority thereafter was to get a billing system operational as this had a long lead time and was essential from a marketing point of view as Newtelco wished to use innovative services as a differentiating point in the market.

Both BT and TD were invited to provide a solution and the board accepted a recommendation from management in July to adopt a solution from Kenan Systems to be integrated by BT Syntegra.

One of the major challenges over the summer for the IT team was to install a new LAN and PABX in the new building in Rumlang, Zurich (see Figure 9), when the company moved from the former BT Switzerland offices in Dubendorf, Zurich (see Figure 10). This was accomplished within very tight timescales and was completed in a quality way, on time and on budget.

Shortly after, Clarify was installed as an order capture and fault

Figure 9—New building in Rumlang



Figure 10—Old building in Dubendorf



management system, as well as the administrative data network (ADN).

At the time of writing, Command 5000, a customer and network inventory tool, a system for computer-telephony integration (CTI) for the customer call centre and a sales and marketing tool are being evaluated.

Products and Services Launch

As mentioned above, the existing WANDA network had some third-party data customers working on it already when it was subsumed into Newtelco. Some 300 customers from the former BT Switzerland were using the Concert CFRS, CPS, CMBS and CVNS services.

To this was added an Internet service provider (ISP) carrier service in July 1997, and domestic data and international single-end voice were added in October.

At the time of writing, installation and testing continue for the full launch on 1 January 1998 of domestic voice products. This will include indirect access from Swisscom using a carrier access code (CAC) and carrier identification code (CIC), international and long-distance voice for directly and indirectly connected customers and an inbound voice service accessed by 0800 numbers.

Further products and enhancements are in the planning and feasibility stage but much depends on the negotiations and action taken to gain a significant reduction in the interconnect rate with Swisscom, which at present is unacceptable at three to four times the European benchmarks.

The Mobile Project

Early in 1998 two mobile licences are to be awarded in addition to the Natel C (analogue) and Natel D (digital) presently held by Swisscom. This is one of the few remaining attractive licences in Europe, as the

penetration of mobile in Switzerland lags the rest of Western Europe, the average international direct dial calling rate is high and the Swiss have a good spending power.

Newtelco will bid for this licence and is adopting an approach which will merge the network and customer offerings with the fixed network. This project was started by a separate project team to allow the people involved in the launch of the fixed network to concentrate on this task. However, this work is now being handed over to the management to develop actively fixed/mobile integration.

A significant amount of risk capital has to be put into this project even before the licence bid but the overall investment is likely to exceed 600 MCHF (£260M).

The Access Challenge

The next big challenge for Newtelco planners is to arrive at a sensible long-term strategy to enable cost-

effective local access to customers. As mentioned above, the present interconnect rates on offer from Swisscom are extremely high. This would tend to encourage investment in either copper or fibre by Newtelco but huge amounts of capital are required to get any reasonable amount of geographical coverage. Swiss law also makes digging extremely costly, and once a highway is opened it cannot be dug up again for two to three years.

Other companies such as MFS, COLT, Cablecom and even some of the city municipalities themselves have fibre access either planned or in the ground so there is the option to reach some form of interconnect agreement with one or more of these companies.

Newtelco has employed consultants to address this problem who have found that, with a very few exceptions, it does not pay Newtelco to build its own local infrastructure.

However, as other technologies mature, such as point-to-multipoint radio, DECT and electronics on the end of copper, this picture is likely to change. Swisscom offers customers 'dark copper' and, with wideband modems at each end, this can provide a good solution.

Personal Perspective

To have the chance to plan, design and implement a network from scratch is a wonderful opportunity. However, it is not for the faint-hearted. At first there were a thousand and one things to do immediately, but no one to do them. A sense of urgency was all pervasive, but the number of hours in a day was unfortunately limited.

As a new company, there were many salesmen and consultants

To have the chance to plan, design and implement a network from scratch is a wonderful opportunity.

knocking at the door trying to sell their own particular product or skills, and seeing them all absorbed a lot of management time.

As people were recruited, gradually order was made from chaos and the rudiments of a quality management system began to take shape. People began to understand the cultural differences that are present and the subtleties so easy to miss in a non-native language. The once strange surroundings became familiar and team spirit began to thrive.

Apart from the new technical knowledge gained by dealing with leading-edge technology and products, the experience gained from a new joint venture brings a good deal of self-development. Learning about

foreign laws, customs, language and life style was, and still is, fascinating. It is to be recommended to any BT person who wishes a chance to exhibit some flair in their chosen field of work and to gain some valuable experience in terms of future career.

Conclusions

This article has tried to convey what is involved when embarking upon launching a new joint venture. The scope of the jobs is much wider than perhaps is normal in a large organisation such as BT, but the personal rewards are great. Mountains can be moved quickly by good team work and the sense of achievement is high.

Postscript

Since the article was written, KomKom has allocated Newtelco the carrier selection codes of 10707, 10701 and 10771. These are equivalent to the 131 and 132 access codes for Mercury in the UK. The last piece of optical-fibre cable needed for opening the network to public service was accepted from SBB on 16 December 1997 and the remaining problem with the billing system interfaces was cleared by Siemens on 30 December 1997. The interconnect links were put into permanent service by Swisscom on 31 December 1997 enabling the network to be opened for service on voice traffic a day ahead of the 1 January 1998 deadline.

Acknowledgements

Many thanks are due to the help from BT and TD staff on secondment to Newtelco and from the parent companies' central engineering divisions. Much of the ability to be first to market was due to previous work carried out by Bernard Mallinder's staff in VIAG and John Skinner's people in Telfort.

Biography



Murray Trantor
Newtelco AG

Since March 1997 Murray Trantor has been on secondment from BT to Newtelco AG, the BT joint venture in Switzerland. He is presently the Chief Operations Officer. Prior to this he worked for BT and its predecessors since leaving school in the mid-1960s and has held a number of senior positions in business systems, operator services and network planning and operations fields in Scotland and the North East of England. His most recent job was Customer and Networks Controls Manager in the CS&NO North East Zone, within Networks and Systems Division. He is a Fellow of the Institution of Electrical Engineers and has been an active member for many years. He is also a Member of the Institute of Management. A keen supporter of IBTE, he has held a number of committee positions and has served on the Board of Editors of its *Journal* for many years. Married, he lives with his youngest two daughters in Wallisellen, a suburb of Zurich.

Ian Cameron and David Patrick

BT WebWorld™—Web Hosting for Every Business Need

BT WebWorld is a new service from BT, launched in October 1996, that delivers a range of Internet Web hosting services based on centrally-managed Web server farms. The service offers customers cost-effective, powerful, flexible and robust Web server capability and associated support services, with BT handling the complex technical and operational issues. This article outlines the background to BT's decision to offer Web hosting services and describes its service architecture, and the features and benefits offered by the product. Finally it takes a look into future service possibilities in this exciting and dynamic market.

Introduction

Once the province of the academic and military communities, the 'Net' is fast becoming a central part of many businesses' communications strategy, and increasingly touches the lives of ordinary people. The World Wide Web, or simply the 'Web', with its user-friendly interface, has played the largest part in the popularisation of the Net. The Web offers users a series of interactive pages, which can include text, graphics, sounds, images and video. These pages are linked together in a seamless way, allowing users effortlessly to browse and select information or services at the click of a mouse. The tool used to access Web information, known as a *browser*, gathers information from one of the millions of Web pages on the Net. It automatically switches between pages and sites to gather the next page of information.

As the popularity of the Web has grown, so has the number of businesses wishing to seize the potential market opportunities that it offers.

Business Use of the Web

As can be seen from Figure 1, the number of public Web sites is growing rapidly, and business users account for the majority of that growth. At the start of 1998, the Netcraft Web Server Survey identified over 1.8 million public Web sites with unique domain names, which is nearly three times more than was recorded in the same survey a year earlier. The actual number of public Web sites (many of which do not have their own domain name; for example, <http://www.btwebworld.com/mycompany>) could be 10 times as large. Add in the private Web sites operating within a company's Intranet (for example, BT has over 1000 such internal sites) and it is clear that the Web is set to have a revolutionary impact on business communications.

Why are businesses using the Web?

Successful businesses regard the Web not as a *network service*, but as a

Figure 1—Growth of the Web (source: Netcraft Web Server Survey, www.netcraft.co.uk)





Figure 2—Business use of the Internet is becoming increasingly sophisticated



market. Buyers and sellers, competitors and partners, products and services, revenue and payment—these are the core elements of the Web market, just like any traditional market. Unlike traditional markets, however, there are very few barriers to entry. It is extremely fast moving and the rewards for the first movers can be great. The Web is being used for communication, research, marketing, customer support, selling and collaborative working. In many cases, there is a natural progression from basic communication and information access towards increasingly sophisticated uses of the Web (Figure 2). Many businesses achieve payback within months owing to cost savings, attraction of new business, speed to market or better market intelligence.

Typical uses of a Web site are wide and varied. Most people immediately think of advertising and selling. But companies are finding new applications to exploit the interactivity, availability and multimedia capabilities of the Web all the time. Some of the most common activities include:

- **Customer services** A Web page can be used to offer 24-hour customer support, seven days per week. The information given can be far more detailed than is possible by post or even on the telephone. Examples include the Railtrack online journey planner and BT's 'Friends and Family' online registration.
- **Complementary marketing sites** These integrate closely into a company's existing marketing activities; for example, Vidal Sassoon or the Year of Engineering Success.
- **News services** These offer instant and up-to-date information around the clock, one example being BT Communication Store. Partnership announcements, and product launches and offers can be covered as well as information resources.

- **Product distribution** This is popular with software companies, such as Microsoft and Netscape, which distribute free trial versions of their products.

- **Micro-marketing** Businesses which have specialist products can use a Web site in conjunction with Internet and other advertising to reach a global market of specialist buyers in the most cost-effective way. There are Web sites specialising in everything from cattle genetics to marine fuel risk evaluation.

- **Web-zines** Web-based newspapers and magazines (such as the BT Businessconnections e-zine) can provide a more rounded view of corporate activities, complement a printed edition, or are just published for their own sake.

- **Complementary telemarketing** Businesses can use a Web site to complement 0800 reception desks. This means that fewer people are needed to man the telephones, and that when customers do call in they are more likely to buy.

- **Market research** Web sites can be used to obtain feedback from customers, suppliers or even employees.

- **Collaboration** Teams that are based remotely can use a Web page to interact regardless of time zones or locations, and are able to meet tighter deadlines.

- **On-line product catalogues** Retail, mail order, and telemarketing-based companies can use Web sites for on-line product catalogues. With the introduction of multimedia on the Web, product demonstrations are also possible. Examples include QualSoft software and Eurovoice.

- **On-line purchasing** It is now possible to sell products and services over the Web using credit card payment, electronic cash and other payment mechanisms. Examples include the BT Shop and Tesco.

External Web hosting

As the Web increasingly becomes a core element of business strategy, so the task of 'hosting' Web content has become mission critical. Few companies, however, have the resources, time or money to build their Web site entirely in-house. Indeed, independent research by Forrester Research Inc. suggests that the cost of setting up an internal Web site can be US\$221 000 in the first year. For this reason, many businesses choose to outsource their Web hosting to network operators, Internet service providers and equipment vendors, which according to Forrester can slash costs by 80%. Although such cost savings are an important factor, it is not the primary reason for outsourcing Web hosting. Forrester's research shows that the need for external labour and knowledge, increased security, increased speed to market and better service and support were all seen as more important in the decision to outsource.

BT WebWorld

Recognising the clear customer demand for Web hosting and related services, BT Internet and Multimedia Services took the decision to offer a service which would revolutionise the implementation and growth of World Wide Web services. Launched in October 1996 and known as *BT WebWorld™*, the service combines BT's extensive communications network and support infrastructure with world-class Web-server software and computing platforms. By mounting Web applications in a secure, networked



Figure 3—BT WebWorld servers in the Cardiff Stadium House server farm

environment developed by BT, customers are free to concentrate on supporting their own business objectives, leaving BT to take on the complex technical and operational activities on their behalf. The service is designed to be as simple or as complex as the customer wants. BT can arrange for complete Web-site design and implementation, using a team of accredited *Web partners*. Alternatively, customers can make use of free tools and online support to build their site themselves. For the more advanced user, the service includes a range of sophisticated features and a built-in library of functions.

Using BT WebWorld, businesses can offer a fast and reliable service to their customers at a fraction of the cost of hosting their own. Customers can have a Web site up and running within an hour, and with competitive monthly costs, few businesses will find difficulty in entering the market.

The product group comprises five main offerings:

- **Publisher** Shared hosting service which offers an inexpensive entry into the Web publishing arena. Aimed at businesses wishing to establish a simple presence on the Internet with minimum cost and complexity, it features a site wizard tool, closed user groups and secure forms.
- **Publisher Gold** Shared hosting service additionally providing a personalised Web address: `www.your_name.co.uk`, or `www.your_name.com` and e-mail address: `managing.director@your_name.co.uk`
- **Premium publisher** Industrial-strength shared hosting Web




Figure 4—BT WebWorld top-level architecture

publishing capabilities aimed at businesses seeking to achieve a quality presence on the Internet. Built on a high-availability platform, the service delivers defined service level targets with high availability and throughput.

- **MarketPoint** This service offers a range of online credit card purchasing services from BT or third party Web sites. BT is currently trialling a service option which is based on Open Market's OM Transact service (see <http://www.openmarket.com>). Another option, using secure forms, was used by Comic Relief to raise over £30 000.

- **Dedicated server** Stand-alone servers (Figure 3) for customers wishing to provide a customised server application. This service offers a managed server platform to which customers have full system access.

Benefits

Benefits of the BT WebWorld service include:

- simple pricing,
- rapid service deployment,
- rapid expansion,
- defined performance levels,
- reduced set-up costs,
- reduced technical risks,
- 24-hour customer service,

- high bandwidth network connectivity,
- high availability,
- 24-hour operations and maintenance,
- security monitoring,
- service creation tools and online tutorials,
- built-in advanced functionality,
- consulting support from Web partners.

BT WebWorld Architecture

The top level architecture of the BT WebWorld publisher, dedicated and Marketpoint platforms is shown in Figure 4.

Content providers' Web pages are loaded onto Web server hardware running Netscape Enterprise server software and BT application software, which delivers the product features. The BT application software interacts with the underlying BT WebWorld Oracle database which contains content provider details. BT WebWorld publisher customers can register online, in which case credit card authorisation is obtained via Barclays Merchant Services who act as the merchant acquirer for this product. Alternatively customers can be invoiced. Additionally, content providers can opt to commerce-enable their content using the MarketPoint platform (at the left of Figure 4) which facilitates credit card transac-


Figure 5—BT WebWorld publisher architecture

tions over the Internet, otherwise known as *e-commerce*.

The publisher platform

Figure 5 depicts the BT WebWorld publisher architecture. This shared Web hosting platform is based upon Sun Sparc 1000 servers in a high-availability configuration, based in BT's Cardiff Stadium House server farm. The publisher application software consists of BT application software residing upon Netscape Enterprise server software. The Netscape Enterprise server module provides basic World-Wide-Web (WWW) page serving functionality. The BT application software builds upon this module to provide product features and basic customer account management, including order/registration, create/delete/modify customer account, suspend/enable customer account and reset password. Content providers are advised of imminent and actual disc space and traffic quota breaches by e-mail. Additionally, closed user group (CUG) management is provided via a WWW interface to CUG functions of the Netscape Enterprise server module such that content providers can specify which end-users can access their content. A set of popular common gateway interface (CGI) scripts are provided for use by content providers, including a form handler and page counter. Traffic statistics relating to content providers' Web sites are made available and can be viewed online via Analog software packages to show statistics such as total traffic delivered, 'hits' per page and destinations of packets. A search engine provided by the service can be used to index Web sites.

Content providers can opt to have secure socket layer (SSL) security, in which case certification is arranged with Verisign, and digital certificates are installed within the Netscape Enterprise server. Also they can opt for hardware or software virtual hosting for 'vanity' Web sites. RealAudio and RealVideo audio and



video streaming are supported on the server. E-mail forwarding can be configured using Nexor software and domain name service (DNS) configuration whereby e-mail associated with a content provider's domain name (for example, sales@companyx.com) can be redirected to a remote mailbox.

For billing purposes, invoicing is implemented via the BT telephony billing and customer service system (CSS). Additionally, the latest release of the BT application software interacts with the BT WebWorld customer database via an application programming interface (API) to facilitate credit card billing and on-line registration. Software processes accompanying the database facilitate credit card billing and reporting. A credit card authorisation module performs an authorisation with the acquirer (Barclays Merchant Services). Subsequently, settlement files are sent via BT's internal networks to Barclays Merchant Services.

Dedicated server platforms

Dedicated servers are particularly appropriate for high-value content providers who do not 'fit' well on BT WebWorld publisher for technical or commercial reasons. This can happen when they do not wish to share a server with other customers, have requirements for extensive CGI scripting or database support, have

high volume traffic requirements, or demand NT platforms. In these cases the dedicated server product provides them with high performance Web server infrastructure tailored to meet their needs.

Customers of the dedicated server product typically have the most challenging requirements for Web hosting. Dedicated servers are usually located in Telehouse (the site of the UK's primary Internet exchange) so as to obtain optimal network connectivity. 'Standard build' UNIX and NT application software are used to provide a basic building block platform. Customers then select which Web server, database and other applications software are to be loaded onto the core platform.

MarketPoint platform

BT is currently trialling the BT WebWorld MarketPoint service (Figure 6), which enables buyers and merchants to buy and sell securely over the Internet. Buyers can use any Visa or Mastercard credit or debit card for purchases, though additional payment mechanisms such as e-cash, procurement cards and secure electronic transaction (SET) are planned for the future. Merchants are responsible for managing and hosting their Web content. BT's central transaction service captures the customer's details, authenticates registered users, calculates the total purchase price and


Figure 6—BT WebWorld MarketPoint service

carries out on-line credit card authorisation and settlement.

Buyers and merchants can view order status and account information at all times using on-line 'smart statements'. In addition, buyers can register online to speed up the buying process for all merchants connected to the BT WebWorld MarketPoint service. Registered buyers need only enter an account name and password at the time of purchase rather than full name, address and credit card details. The service supports the purchase of soft goods (for example, information or software which can be downloaded over the Internet), hard goods (goods or services which are delivered through other channels; for example, by post), and subscriptions (soft goods to which the buyer can subscribe for a period of time; for example, an on-line publication).

Not all products can be sold effectively over the Internet. Those which are most successful tend to be attractive to the Internet audience (predominantly male, technically aware, between 20 and 35 years old). Figure 7 summarises some of the products which are currently selling well in the business and home buyer markets.

For the BT WebWorld service, *commerce enablement* is implemented using Open Market Inc.'s 'digital offer' technology as shown in Figure 8. This technology intercepts buyers (end-users) whenever they attempt to download commerce-enabled content from the WWW



server, and requires them to identify themselves and purchase the content before they can access it. Buyers are redirected to the BT transaction service for online registration or card validation, and the card details are passed to the card processor for authorisation and settlement. Once the transaction is authorised, the buyer is issued with a 'digital receipt' which gives them access to the content on the merchant's Web site.

A similar model is used for hard goods purchase and subscriptions. Funds are transferred directly from the buyer's credit/debit card account to the merchant using the merchant's own card processor merchant-ID.

The use of a central *transaction server* model offers significant improvements in security over in-house solutions. All information transmitted between the buyer and the system (including credit card numbers) is encrypted using the secure sockets layer (SSL) protocol. Credit card details are encrypted on the transaction servers, and are kept encrypted until they are sent to the credit card company over a

private network. Indeed, merchants never see the credit card numbers, and buyers can only update their credit card numbers—they cannot view their numbers. The user name and password of registered buyers are supplemented by a *custom challenge* system. This provides a secondary level of security for transactions above pre-specified limits. Buyers can also set financial limits for purchases made using the system; for example, only allowing purchases up to £50 in value unless a custom challenge is answered. As only BT handles secure customer data, buyers have no need to worry whether each individual merchant is genuine or has correctly implemented their security procedures.

Content authoring

For simple hypertext markup language (HTML) content authoring, content providers are able to use the Forman Internet Creator Site Wizard supplied by BT. This is a simple Windows-based tool which allows customers to create a complete site without any knowledge of HTML or other Web authoring techniques. The tool features one-button publish capability making it possible to develop and publish a complete Web site in under an hour. Many customers use commercially available Web authoring tools to create more advanced sites.


The Future

It is challenging to predict accurately the long-term future of the Internet. The main technology trends that will

Figure 7—Successful products for Internet commerce

	BUSINESS BUYER	HOME BUYER
HARD GOODS	COMPUTING OFFICE STATIONERY PAPER REPORTS COMPONENTS/ REORDERS	COMPUTING BOOKS MUSIC (TAPE/CDs)
	JOURNAL SUBSCRIPTION INFORMATION (FINANCE, LEGAL, REPORTS) SOFTWARE AGENCY SERVICES	ENTERTAINMENT, MUSIC EDUCATION, HOBBIES SOFTWARE, AGENCY SERVICES INTERMEDIARIES

Figure 8—Typical BT WebWorld MarketPoint soft goods purchase flow



drive new features in BT WebWorld over the next few years are likely to be technology/standards advances, telephony, real-time services and electronic commerce. Among technology/standards advances we expect to see intelligent advertising and customer profiling, increased HTML standardisation and a proliferation of companies becoming certification agencies. In the areas of telephony and real-time services we anticipate the incorporation of computer telephony integration with Web pages, multi-party Web conferencing, widespread real-time telephony support and multicasting. As end-user bandwidth increases, the boundaries between television, radio, telephony and Internet services will become increasingly blurred. The emergence of high-quality real-time video and audio streaming over high-bandwidth multicast networks will create new personalised entertainment and information markets over the Internet. In the area of electronic commerce, WWW implementations of electronic data interchange applications, the introduction of SET technology and mature micro-payment schemes may make the Internet the primary electronic payment infrastructure.

Overall, the Web has moved from academic novelty to business-critical service in a little over two years. This is a revolution which seems certain to accelerate, creating significant opportunities for the users and equally significant challenges to the developers.

References

Information about SET: <http://www.visa.com/cgi-bin/vee/sf/set/intro.html?2+0>

Information about CGI: <http://www3.pair.com/webthing/docs/cgi/faqs/cgifaq.shtml>

Information about SSL: http://home.netscape.com/newsref/std/SSL_old.html


Biographies



Ian Cameron
BT Internet and
Multimedia Services

Ian Cameron joined BT in 1981 on the University Studentship Scheme. After graduating with a degree in Electrical Engineering from Imperial College in 1985 he joined BT's Research Department where he developed interactive speech services, including telephone competitions and telephone banking. In 1991 he joined BT's Product Management division where he led the product launch of new calling card services such as VisaPhone. In November 1993 he became manager for BT's video-on-demand trial and more recently has led product development of international broadband services and groupware services. He is now product manager for Web hosting services within BT.

ian.cameron@btwebworld.com



David Patrick
BT Networks and
Services

David Patrick joined BT in 1992 after graduating with a Ph.D. in Physics from Imperial College. He initially joined BT's Research Department where he published a number of papers in the field of ultrafast fibre optics. Subsequently he joined BT's Network Studies unit where he pursued the topic of Internet evolution. He currently works within BT in the role of Internet and online manager, with technical design authority responsibility for Web hosting and related content services within BT.

patrickdm@btlip02.bt.co.uk

Games People Play—The Wireplay Experience

People have played games for thousands of years using boards, cards and dice, and more recently on computers and over networks. Although there is a bewildering number of games, there are only a few basic types. Wireplay is the games network from BT that provides smooth play with all types of networked computer games. This article describes Wireplay from a user's point of view.

Introduction

People have devised and played games for thousands of years. Despite appearing to be a frivolous waste of time and resource, playing games is a powerful means of gaining and maintaining valuable skills by both children and adults.

For instance, the ancient game of go (Figure 1) encourages the development of strategic thinking and mental agility. Other games develop physical skills such as hand-eye coordination, which is why English yeomen played regular games of darts to maintain their accuracy in archery.

Games can be classified by the playing equipment they use; for example, board games, card games, dice games, computer games. Although there is a bewildering number of different games in the world, there are only a few basic types, which are described below.

Turn-Based Games of Skill and Chance

These are the classic games where players take turns to make a move.

The game progresses through a series of states and the players attempt to reach a state which meets certain criteria in order to win. Each game has a particular mix of skill and chance.

Games of pure skill such as chess, involve chance only in choosing who goes first. Chess (Figure 2) is a game with an enormous number of states, making it impossible for the human mind (or computer) to look ahead and evaluate all the possible outcomes of a move. Successful play requires great skill and concentration, and challenges the finest intellects. At the other extreme, noughts-and-crosses is a simple game where most people with a little practice can guarantee at worst a stalemate.

Some games such as rolling dice are based on pure chance. They are often livened up by betting, and successful play can include highly skilful elements of bluffing, risk evaluation and applied psychology.

Many games are between these extremes with enough chance to make the outcome unpredictable and enough skill to challenge the players. New games are constantly being devised but

Figure 1—Wireplay go




Figure 2—Wireplay chess



if they get the balance wrong by being too predictable or too complex they do not become favourites.

Before network gaming was available, turn-based games were not a popular subject for computer game development. Programming the computer to play a game like chess was difficult and the computer tended either to be far too good or to make deliberate mistakes, both of which are unsatisfactory for a keen player. However, the number of people connecting their PCs over networks to play games with other people is increasing rapidly and this has renewed interest in turn-based computer games. Turn-based games are well suited to playing over the Internet since any delays in transmission are masked in the passing of play from one player to another, although a more responsive dedicated games network like Wireplay provides smoother play. Wireplay offers chess and a free bridge game for its users and several other titles are due to appear in the near future.

Games of Dexterity

Games of dexterity include most sports where speed and hand-eye coordination are important. They are often timed and may be played in teams.

Platform games (Figure 3) are a type of computer game which test

Figure 3—*Sonic and Knuckles*—a one- or two-player platform game (*Sonic and Knuckles* is a trademark of Sega Enterprises Ltd.)



dexterity and mental agility. In these games, the player controls a character that runs across the screen jumping from platform to platform, avoiding pitfalls and collecting objects to accumulate a high score. Memory is also required to remember the best route through the game.


These games can be enormous fun for one player and have been produced in two-player variants for game consoles where two control pads are used with a shared screen. No networked variants of platform games are known, although dexterity forms an important aspect of other genres.

Role-Playing Games

People of all ages enjoy playing 'let's pretend'. Much of children's play is based on elaborate role-playing, and many adults enjoy games such as costumed re-enactments of historical events and other fantasy games. These games have considerable value in providing an opportunity to try new behaviours and consider other viewpoints.

In computer-based role-playing games, the player either steers a character or effectively becomes a character by adopting a first-person view. A virtual world is presented through which the player navigates to find objects and achieve a goal. The virtual world is populated by

Figure 4—*Duke Nukem 3D*—a first-person shoot-'em-up game (image courtesy of Eidos Interactive)



other characters that are either computer-generated or controlled by other players over the network.

Some of these games simulate violence including the variety based on armed combat known as *shoot-'em-ups* (Figure 2). Probably the world's most popular game in this genre is Quake in which characters with a huge range of weapons fight to capture a flag. Other games (for example, Myst) stress problem solving and are gentler in nature. In each case the world is made as realistic as possible so that players feel that they are 'really there'. Quake also gives players the means to modify the world and their appearance in it.

A special variant of the role-playing game is the text-based adventure game in which players use their imaginations to visualise the action. Wireplay offers its users free software to play MUD II, created by Dr Richard Bartle, one of the originators of this genre. Players find this world very involving and may spend long periods on-line to develop their characters. Wireplay also offers more traditional shoot-'em-up games.

Resource Management Games

These are games in which the player directs resources to achieve some goal such as accumulating wealth or winning a war. The resources appear to be autonomous and their behaviour is not predictable, so that skill is needed in deciding when and what to do with them. Business simulations use this model, but many resource management games are played for pure entertainment. Computer-based resource management games are sometimes known as *god games*. Some very popular titles have been produced in this genre including Populous and Theme Park, both for one player.

Networked resource games allow more than one player to compete. In Warcraft 2, which can be networked via Wireplay, either orcs or humans are controlled. Each player directs



Figure 5—*Fragile Allegiance*. Space resource game with battle mode (image courtesy Gremlin Interactive)

the training of his or her forces and creates an army to fight the opposing forces. The challenge for players is to balance tactics and strategy; for example, deciding when to switch from training and organising to fighting. Other titles supported by Wireplay include *Command and Conquer—Red Alert* and *Fragile Allegiance* (Figure 5) and many others are in preparation.

Simulations

Simulations approximate the experience of flying fixed and rotary winged aircraft (Figure 6), cars, tanks and many other vehicles. The best come very close to real life with realistic displays and instrumentation taking hours to master. Once restricted to training professional pilots, sophisticated simulations are now in widespread use by computer gamers.

Simulations are the only genre uniquely enabled by computer, all the

others could be played with dice and a rulebook. Simulation games have a long history from the early flight simulators to the latest EF2000 simulator offered on Wireplay, which is based on the simulator used to develop the real aircraft. Other vehicle simulations are available, from cars (*Big Red Racing*) to helicopters (*Commanche3*). Networked multiplayer simulations allow players to race or fight against each other.

The Need for Speed

When two players meet in a shoot-em-up game they need to see each other immediately and there must be no doubt as to who fired first! The same requirement is found in other role playing, simulations and dexterity games.

Even in turn-based games, a small delay repeated time after time will irritate many players and if delays approach a second or more they begin

to be unacceptable. Play is always smoother and more fluid if delays are kept imperceptible so that the technology fades into the background and players become fully immersed in the game. The same applies to resource management games and, in multiplayer versions, the resources must interact smoothly so that they respond immediately to an opponent's action.

Consistency of response time is also important so that players can accurately predict the outcome of their actions and to prevent systematic bias in response time making a game unfair.

Networking games by means of a local area network with sufficient bandwidth allows response times of a few milliseconds. However, players would have to travel to the network in order to get this kind of experience. Networking over a dial-up connection results in greater delays, and access via the Internet introduces a much greater degree of variation in response time. Players can tolerate up to about 400 ms latency if the game include a 'predictive engine'. This calculates what is likely to happen next and uses it to fill in the gaps between updates. When response times exceed 400 ms games become jerky, as they make noticeable corrections when the updates arrive.

Wireplay achieves less than 150 ms latency with a high level of consistency, resulting in an excellent player experience with most games. Further improvements are planned to reduce the need for predictive technology and improve the performance with games that have not been written especially for on-line play. The technical detail of how Wireplay achieves its high level of responsiveness is described elsewhere†.

Figure 6—*Flying Corps GOLD*. World War 1 flying aces dogfight (image courtesy of Empire Interactive)



† MILLAR, WILLIAM; ASHBY, PHIL; DUFFY, MARTIN; UPTON, COLIN; and WELSBY, KEVIN. A High-Performance Network for Gaming—Wireplay. *Br. Telecommun. Eng.*, 16, Jan. 1998 (this issue).

Figure 7—Wireplay Noticeboard


The Wireplay Experience

A few games (for example, Quake and MUD) have become established as classics. However, most games have a much shorter active life. Multiplayer game playing extends this life, but the large number of publishers vying for custom and constant advances in technology encourage the proliferation of new games.

Wireplay has therefore been engineered as a service for people to meet and play a large number of different games. A sophisticated user interface makes it easy to go on-line and find out what games are currently being played. This includes filters to select the desired playing standard to avoid challenging a badly-matched opponent. The interface will be extended to make it easy to navigate the large number of games that will become available as the service evolves worldwide.

Users can obtain Wireplay software from cover-mounted CDs on popular PC magazines, or direct from the Wireplay Web site—<http://www.wireplay.com>. The software is self-installing and comes pre-configured to connect to the service. Once installed and launched, a set-up screen allows users to check their modem selection and add the games they own using a built-in list of supported games. Some free games are included to start them off.

When ready, the user is prompted to enter a nickname and to connect to the service. Once on-line, the service checks the uniqueness of the nickname, asks for a password, registers the user's games and then displays terms and conditions. If the user accepts these, he or she is connected to the service and sees Wireplay News, a compilation of information about the latest games, competitions and news about the games industry. They can then move to the Noticeboard (Figure 7), the heart of the Wireplay system. This shows how many people are on-line, playing games and how many games proposals there are.



To play a game, the user either selects an existing proposal or creates a new one. Players enter a 'Locker Room' in order to exchange text messages about setting up the game and when everyone is ready, Wireplay starts the game on their PCs and manages all messaging between them. When the game finishes, Wireplay registers the result and returns the players to the Noticeboard.


The objective of the Wireplay service is to provide a complete surround for game playing and develop this to be what players want

and need. The service has numerous features in addition to the basic description provided here and is under rapid development.

Future Developments

Wireplay has been an operational service since June 1996 with full service from October 1996. The latest version includes Windows 95 software with more functionality and even better communications. The rate of innovation demanded by the games industry makes it very difficult to predict future developments.

Figure 8—Wireplay Locker Room



Wireplay is being adopted by partners around the globe, starting with MCI in the US and Telstra in Australia. The Wireplay service is increasingly being seen by games publishers as a serious and ubiquitous platform for multiplayer game playing, making them more likely to write games to take advantage of its special features.

Work continues to achieve further improvements in system response times and to minimise the effect of the traffic imbalance that is inherent in on-line game playing. In an N player game, there is $N-1$ times more traffic incoming from the other players to a player's PC than outgoing from that PC to the network. This can cause problems with Internet connections owing to delayed messages and the message order becoming garbled. In Wireplay this problem is minimised in three ways. For Wireplay games, processes in the Wireplay server aggregate game information so that information from all players is packaged up into a single message which is routed to everyone at once. For non-Wireplay games, software specially written by the game publishers to minimise traffic between PCs is provided and more is being added. Finally Wireplay uses Internet technology in a managed intranet with careful traffic control to prevent bottlenecks and delays.

Concluding Comments

Networking computer games widens the range of opponents available to a player; overcomes the limitations of geography and adds a social dimension to what can otherwise be a solitary pastime. The Wireplay team is working with the games industry to ensure that the service provides the best possible performance with the widest choice of games. For more information about Wireplay, visit the Web site at <http://www.wireplay.com>.


[Note: All images are from public Web sites and are current at the time of writing.]

Biographies



Kelvin Lawrance
BT Networks and Systems

Kelvin Lawrance joined BT in 1987 after graduating from the University of East Anglia with a B.Sc. in Electronic Systems Engineering. After several years working in areas including videotex applications, directory enquiries services and office automation systems, Kelvin first became involved with the games sector during the interactive television trials, while working on a service delivering content to games consoles. He has been involved with Wireplay since its inception and early design, and has fulfilled a multitude of roles during the service's evolution. He currently leads the teams providing support to both games developers and licensees.



Dave Cleland
BT Networks and Systems

Dave Cleland joined BT Laboratories in 1989 having graduated from the University of Bath with a first-class degree in Electrical and Electronic Engineering. After three years of developing high-speed transatlantic optical-fibre submarine systems, he went into software systems development with a special interest in object-oriented software techniques. During this period he designed and developed areas of BT's CallMinder and MeterLink projects. Since June 1996 he has been part of the Wireplay team, where his role has been working with UK PC game development companies to encourage and help them to create games that will operate smoothly over BT's network.



Bryan Johnson
BT Networks and Systems

Bryan Johnson joined BT Computing Information Services in 1988 as an analyst/programmer working on MACE—a mainframe information system. In 1991, he became a systems engineering consultant in BT Information System Development division supporting the adoption of formal system development methods throughout BT. He joined SID as a test team leader in 1993 working on customer service system (CSS) network systems. His responsibilities included the management of the testing and integration of CSS network subsystems and the operations and maintenance centres. He joined the Wireplay project in 1996 to take responsibility for a team testing and integrating PC computer games software with the Wireplay software and coordinating the release of those games onto the operational service.



Kevin Welsby
BT Networks and Systems

Kevin Welsby leads the Wireplay development team. He has worked with BT since GPO days, starting as an apprentice. Kevin graduated from Aston University in 1978 and moved to BT Laboratories at Martlesham Heath where he first worked in the acoustic laboratories. He delivered several customer products including telephones and payphones before moving into IT. As systems architect he developed designs for CLASS and telemetry projects before moving back into Products and Systems Engineering to deliver the solutions. He took over the Wireplay project in late-1996 in order to build a high-performance team which would turn Wireplay into an operational service.

A High-Performance Network for Gaming—Wireplay

BT's Wireplay has been designed to provide a unique game playing experience for customers. Every aspect of the system has been optimised for fast and efficient performance, resulting in a best-of-breed games network that is being exported worldwide.

Introduction

The Wireplay service from BT enables PC owners to meet, chat and play interactive games over the telephone network. Customers obtain free Wireplay software, often bundled with games, from a number of distribution channels such as the CDs mounted on the covers of PC magazines. The customer installs the software and is able to play immediately without needing to subscribe to a service provider: Connection is via a modem to an 0894 number, a premium number range for gaming with an attractive tariff structure for gaming. This simplifies billing for the customer and BT, although more conventional billing is also proposed to allow credit card subscriptions. Wireplay is described from the user's perspective elsewhere†.

This service sits at the juncture of communications, computing and information technology, and as such presents significant challenges to the design, development and operational teams in order to maintain a good quality of service to the end customer. In particular and uniquely, Wireplay is a dedicated games service with a proprietary communications protocol and optimally configured servers and network for fast and efficient game playing. On-line games services based on standard Internet technology cannot achieve the same level of performance or consistency of performance under load.


† LAWRENCE, KELVIN; CLELAND, DAVE; JOHNSON, BRYAN; and WELSBY, KEVIN. Games People Play—The Wireplay Experience. *Br. Telecommun. Eng.*, Jan. 1998, 16 (this issue).

Software and Hardware Architecture

The architecture of the current implementation of Wireplay is shown in Figure 1.

Communication between the user and Wireplay server is via client software that also handles the user interface and management of games for the service. A proprietary protocol, PFS, is used to give the lowest possible latency connection to the Wireplay terminal server. PFS is the subject of a patent application. Once connected to the Wireplay server all signals are carried in a high-speed local area network (LAN) environ-

Figure 1—Major system components for the Wireplay service



ment so standard Internet (TCP/IP) protocols can be used from this point. The function of the terminal server (apart from network termination) is to convert from PFS to TCP/IP protocols.

On first connection, and as needed during operation of the service, the router server acts as a proxy, routing data to the Noticeboard which is the centre for servicing the users' requirements, including registration into the service, authentication of passwords and matching sets of users to play games. Users register by nickname and although further customer information is required by the system, this is not made available to other users. Customer details are stored in a database which allows powerful search and query functions for both customer service and service management activities.

Once a group of users is matched for playing a game, the router server is able to route game data unaided. The proxy session is made idle but maintained for use after the game, and the router server passes data from client to client as required by the game. In simple implementations the router just provides client-to-client messaging, but for games which have been specially written to use the feature it is able to aggregate data from several players and send this as one packet. This reduces overheads and handshaking to achieve further improvements in responsiveness. Alternatively, game traffic may be routed to a game server that runs software dedicated to a particular game. This server can process game data so as to minimise the traffic to and from each player.

On completion of the game, the proxy session to the Noticeboard is re-opened and scoring information may be passed to the database for ranking players, competitions and so on. The database is also used to give management information for the Wireplay team, via the service management system. Reports on service statistics against a number of criteria such as number of games played, players per hour etc. can be

generated and made available for further consumption by marketing and development teams. Reports may be routinely available or tailored for specific purposes.

Platform Implementation

Currently the Wireplay service is implemented on general-purpose Unix computers. These give adequate performance when configured as described and minimise the capital risk to BT as they can be re-used for other products.

The Noticeboard and database reside on one server optimised for database queries and customer interaction, and this will probably always be hosted on a general-purpose computer. A relatively small system is adequate for these non time-critical applications, although more capacity can be added if required.

The router is implemented on a second computer with a highly optimised and near real-time Unix operating system, pushing this technology to the limits. A number of processors can be incorporated in the router and the router replicated for increased capacity, which makes the solution very scalable. This time-critical part of the system could be replaced with a purpose-built real-time system if required, which would reduce the cost per channel but have relatively little effect on performance as this is already optimised.

Access and Network Communications


The ability to control latency (crudely the time for data to leave a game on one player's PC and be delivered to another) is crucial in providing a good game-playing experience for the customer. Even if the game includes predictive methods to disguise delays and missing data, slow networks will cause jerkiness while the game overcomes deviations from prediction or struggles to catch up with the true position in the game of all the players. Under these conditions, some games

may even halt altogether for several seconds before play resumes.

Internet technology is good at reliably delivering sizeable packages of data and can do so at reasonable data rates if the network links are good. However, it achieves this by halting and queuing packets. The resultant delays are not noticed in typical applications such as Web browsing, in which users may download several images of a few kilobytes each on a page and are generally happy if the images download at a reasonable average rate. On the other hand, game data is typically in much smaller packets, down to a few bytes, and must be delivered as quickly as possible. In these circumstances a high data rate achieved by accumulating and sending a number of packets every few seconds is unacceptable.

Wireplay overcomes these problems by using highly optimised dedicated network components in the service platform. The access components are shown in Figure 2.

Figure 2—Wireplay access and network components



Wireplay has its own dedicated point of presence built from Ascend Max 4000 modem banks. Each customer's call is delivered to the modem banks over the public switched network using integrated services digital network (ISDN) bearers, guaranteeing the lowest possible delay from the customers' modems to Wireplay. The modem banks then act as terminal servers receiving data using Wireplay's proprietary PFS protocol and packetising this for transmission on the high-speed Wireplay backbone using standard TCP/IP protocols. The proprietary protocol is used to achieve the highest possible performance over the low-bandwidth link from home to service while standard protocols are used in the high-speed environment where they do not affect performance.

Currently a 10 Mbit/s Ethernet LAN is adequate for Wireplay, although work is already under way to increase this to meet expected demand. The LAN cannot be used in a standard way because the delays and variability of latency incurred by waiting to fill up the IP packet before transmission would be unacceptable. The software in the modem bank has been enhanced to flush the modem on receipt of the 'end of packet' marker in the PFS protocol. The software also includes a timer that is set at the start of each packet to a period greater than the maximum expected packet duration. The timer counts down as the packet arrives and flushes the packet if it reaches zero before an end-of-packet marker is recognised. This prevents the system coming to a halt if the marker is corrupted. Careful optimisation of the timer setting and packet size gives Wireplay its good performance. If the timer is set too low, packets are flushed before they are complete and the data becomes fragmented. Although the protocol reassembles the packets later, this causes delays and jitter that affect the game-playing experience.

Wireplay does not insist on particular packet sizes and data rates. However, BT gives a high level of support to games developers who choose to collaborate with BT in integrating multiplayer communications. Game requirements vary widely from burst transmission of hundreds of bytes or more during games set-up to very small packets of typically 14 bytes during play. Many games have been developed for LAN play and adapted, with varying degrees of success, for play over the telephone network using modems. The networking has been developed pragmatically for good performance with numerous games, and work continues with key games developers to establish the ground rules to ensure that their titles offer optimal networked performance with Wireplay.

Testing

Many claims are made for the performance of competing games services, typically based on 'ping time', the time for a packet to be reflected back to the sender. Times are quoted but not under load and to be fair this is not easy to do as conducting a realistic load test is not a trivial exercise. Load testing with real people is difficult and so a number of automated test tools have been made to assess the performance of Wireplay.

Latency tester

The Wireplay latency tester emulates real game play by sending packets across the entire network as if they were from a client PC as shown in Figure 3.

A dummy game, Wplat is run on the service which can be entered as usual by proposing and starting the game from the Noticeboard although this game is not visible to normal customers. A second PC joins the game to run a variant, Wprelect, whose purpose is simply to return packets as received. The Wplat machine can be set to transmit time-stamped packets of virtually any size and continuously monitors the time

taken to receive them back again. The times can be listed or displayed graphically to show immediately the effect of changing system parameters.

Because the data is transmitted to a second machine there can be an assurance that all elements of the Wireplay system are exercised and that the measure is fair. In early development this tool revealed some deep interaction effects which were causing unexpected performance problems that could not have been discovered in any other way.


The latency test can be used in parallel with the load test as a powerful diagnostic exercise and is run continuously on the live service as an early warning system. As a controlled source and sink of packets, it is also the foundation of a number of other tests that examine the performance of the platform at various points.

Load tester

Testing the performance of the Wireplay system raises questions such as:

- What are true game playing conditions? Should attention be concentrated on the most demand-

Figure 3—Latency testing



ing games in terms of the traffic they generate or is it more realistic to take a mix of games and if so what is the right mix?

- Should the service be tested fully loaded with games only or should a mix of other operations be taken such as chat, file download, news browsing etc? If a mix, what is a realistic proportion of each activity?
- Are all players going to drive the system hard, or should a range of human reaction times be assumed?

The answers to these questions are only becoming clear as the service evolves and could not be guessed at the start of the development work.

The ideal load tester would emulate a standard and an extreme mix of operations on the platform and exercise the whole system. This would require as many modems as the intended system capacity and a great deal of computing power to drive all of them with emulated games and other traffic. Realistically this cannot be achieved owing to commercial considerations—the power required and the capital cost would exceed that of the platform itself several times over. This would also lead to an ‘arms race’ in which increasing platform size with demand would require even more load test capacity. Instead a two-level strategy has been adopted. First a load test capacity has been developed which can load up the high-speed Wireplay network and servers. In addition, capacity is sufficient to allow any one of the current modem banks to be fully loaded. In fact, capacity exceeds this to allow for the higher-capacity modem banks expected in the future. The load test system is shown in Figure 4.

The load server runs on a general-purpose UNIX system and emulates users creating games proposals and playing games. It could run on one

large computer (for example, testing the live site using the reference model as load generator), but in general all the large systems are tied up during a load test as they are configured into the system. Instead a number of smaller boxes, known within the Wireplay team as *froggers*, are used in combination to generate the load.

The network load tester consists at present of four low-end PCs driving 16 consumer modems each. This is easily scalable by adding more PCs and modems and is about the lowest cost means of exercising the public network access part of the system. Again, software on these PCs emulates users creating games proposals and playing games.

Performance testing

In combination the above tools give a powerful cost-effective method for examining the performance of Wireplay, in detail and under load. The network access point can be fully loaded and operated into a fully loaded Wireplay server. The latency tester can be used to exam-

ine overall objective performance under full and partial load. Detailed diagnostics and standard LAN analyser tools can be used to determine the cause of any deviations from optimal performance. The whole set-up can reasonably be assembled whenever needed for assessment of new software releases or to evaluate new systems elements.

While this level of testing has proved vital to the engineering success of Wireplay, it is easy to miss the point with objective measurements. At the same time as performing the load tests, it is important to include assessment by a few expert users, as the subjective performance of the service as viewed by the service’s very demanding customers is of paramount importance.


Games Integration and Testing

The games-integration team makes the bridge between Wireplay technology and the needs of game-playing customers. Customers primarily judge the service by how well games play and whether Wireplay can support at least as many players in a game as its competitors.

Wireplay supports a number of game types, both DOS and Windows 95:

- Wireplay multiplayer application programming interfaces (APIs). Wireplay provides software for inclusion in games which incorporates the communications layer for the game developer. The game makes calls on this API and all underlying communications are hidden from the game developer. The API includes added functionality for registering the results of games so that the system is able to keep scores, rank players and provide competitions features.
- IPX networking through the WireIPX™ system. Many games have been written for LAN

Figure 4—Wireplay load tester



operation and, provided the data rate is not excessive, these can be played over modems. Such games are becoming more common as competitor's services have based their operations around them exclusively. However, play can be very variable and these games are assessed very carefully to ensure that the games are playable. A preferred maximum number of players is recommended and an absolute maximum is enabled on the system, for players who can accept minor problems in smoothness of play.

- IP (Internet protocol) networking through the WireIP™ system.
- Microsoft DirectPlay™ version 3.0 or above.
- Special Wireplay-produced drivers for certain popular games, which accelerate performance over the Wireplay network.

Only Wireplay specific games are fully integrated. The others require manual intervention by the customer to various degrees before the game can be started. Often this means only that the customer has to select 'network play' once the game has started, as the software is not able to detect that Wireplay is present as a network layer.

Wireplay API games are developed in collaboration with the game developer and follow more conventional software test procedures. Other games must be tested from a

user's perspective and Wireplay maintains a games testing laboratory equipped with a range of PCs covering the types likely to be in use by customers. These range from machines meeting the Wireplay minimum specification up to the latest MMX equipped PCs with 3D-graphics cards and all the goodies.

Wireplay is not responsible for the content of games and titles will be published with or without BT's blessing. BT is guided by age ratings as published by European Leisure Software Publishers Association (ELSPA) and would not normally enable a title onto the service that is unacceptable to the industry. Game titles are constantly evaluated both before and on publication, and a weekly conference is held to review potential new titles.

Several copies of new titles are required and systematically evaluated before release into Wireplay. Cooperatively developed games are tested with the developer and problems are referred for resolution as found. Titles that are ready developed are tested for compliance with other supported formats. For both types, the number of players that can join a game with faultless play is tested and the maximum that can play without major failures. Any foibles are documented and a frequently asked questions (FAQ) sheet is prepared for inclusion in on-line information services, which helps users get the game going and understand any oddities.

To enable the game onto the service there has to be a databuild on both client and server. Games are enabled onto the server via the service management system, after which the game can be issued within Wireplay for key people, such as BT Helpdesk staff, to try them out and be ready to answer questions. The game is not visible to users until they update the games file on their PC. This is easily done by downloading a small file from the Wireplay News page where the availability of new games is also promoted.

Service Platform Operation


The Wireplay service platform is currently in operation at Bibb Way, Ipswich. Its proximity to the engineering team has helped development of the service in early life but growing capacity demands will dictate a move in the near future.

An operations manager who is responsible for meeting in-service targets for uptime and repair turnaround manages the platform. The service has rapidly achieved very good availability and is now better than 99% available. The service platform is supported by BT's Customer Service Operations (CSO) remotely from the operations centre in Sheffield. Application alarms are generated and collected on the servers before being sent to CSO via standard OSMF interfaces. A customer helpdesk provides direct customer support from Colindale seven days a week from 10.00 hours until 22.00 hours. Second-line support for the application, including server resident games, is given by the Product Support Group within the Systems Integration Department of Systems Engineering.

Also from Systems Integration Department, a team of verification, validation and testing (VV&T) specialists work with the developers from specification to delivery to ensure the quality of the product and that customer problems are minimised. They also maintain configurations of the reference system so that versions of the system both for the UK and overseas can be worked on as required.

All the above have responded magnificently to the demands of Wireplay. The games industry is still immature and fashion-led, which has resulted in very high-speed delivery to market and a great deal of volatility in plans. BT people have moved rapidly and taken responsible risks while maintaining the standards that BT expects.

Figure 5—Wireplay games laboratory




Conclusion

Wireplay is not an especially complex system although it has presented considerable challenges to achieve the desired performance. It operates in a chaotic software environment. The market is also chaotic and the teams have had to achieve a very high rate of delivery, with few months not including some significant deliverable. There have also been cultural barriers to overcome, including working with an industry which is very different from mainstream BT. With Wireplay, BT has shown that it can be part of this exciting world and achieve world-beating performance. The technology is being exported worldwide and a very significant upswing in interest is being experienced from game developers. The game has just started!

Biographies

William Millar
BT Networks and Systems

William Millar joined BT in 1986 from The Queen's University of Belfast, where he was a lecturer in the Electrical Engineering Department. In BT, he worked on speech recognition and neural network algorithm research and development. He then moved on to lead a team responsible for downstream signal processing algorithms to the BT speech applications platform (now the ETL ISAP). Since then he has led a corporate research campaign on multimedia platform futures and has been the delivery manager for some of the BT WebWorld product set. He has been product engineer for BT Wireplay since November 1996, during which time it has been launched in the USA and Australia and its capability has been extended in the UK.



Phil Ashby
BT Networks and Systems

Phil Ashby joined BT in 1985 as a sponsored student. After graduating from York University in 1989 with an M.Eng. (Hons.) degree in Electronic Systems Engineering, he joined the Speech Applications Division (later to be known as Products and Services Engineering) to work on voice-controlled telephone banking systems. After this, he worked on the NBCAS/CallMinder project until its initial launch in 1994. He then joined the Speech Applications Platform (SAP) team, working with large real-time systems until the SAP development was transferred to Ericsson in 1996. Since then, he has been part of the Wireplay development team, involved in the networking and game-integration aspects of the system.



Martin Duffy
BT Networks and Systems

Martin Duffy joined the multimedia applications division of BT in 1988 and developed a number of Microsoft Windows based office automation applications before moving into the area of multimedia kiosks including BARBARA (access to local council services from remote locations), City Media network services (remote access to multimedia databases) and BT Global Challenge (access to the race results system). He currently leads a team developing Windows-based applications for the BT Wireplay service. Martin holds a B.Sc. (Hons.) in Electronics and Computing from the University of Essex.



Kevin Welsby
BT Networks and Systems

Kevin Welsby leads the Wireplay development team. He has worked with BT since GPO days, starting as an apprentice. Kevin graduated from Aston University in 1978 and moved to BT Laboratories at Martlesham Heath where he first worked in the acoustic laboratories. He has delivered several customer products including telephones and payphones before moving into IT. As systems architect he developed designs for CLASS and telemetry projects before moving back into Products and Systems Engineering to deliver the solutions. He took over the Wireplay project in late-1996 in order to build a high-performance team which would turn Wireplay into an operational service.

Steve Buttery, Bob Merrett, Adrian Pote and Andy Sago

Local Area Mobility

BT intends to exploit the current high demand for mobility by creating a new range of products that build upon the best features of existing 'mobile' offerings. Under the name OnePhone, these products will enable business and residential customers to access both digital cellular and digital cordless services via a single 'dual mode' handset. In this way, OnePhone can become an important step towards the universal mobile telecommunications system (UMTS) – Europe's vision of telecommunications in the next millennium.

Introduction

One of the most important current trends in telecommunications is the increasing demand for mobility. Perhaps the most obvious example of this is the mobile telephone market, which is set to more than double in size over the next five years. There are, however, other ways in which customers are getting the mobility that they want, including:

- **Cordless Systems:** By connecting cordless systems to their fixed network connections, customers can have mobile access to the network from their homes or offices. This allows them to have freedom of movement around their premises while retaining the low cost and high quality of a fixed network connection.
- **Personal Mobility:** By subscribing to a personal number service, customers can ensure they receive calls via the most convenient network connection, wherever they are.

Looking at these different ways of achieving mobility, it is clear that they all have their own particular strengths:

- The global system for mobile communications (GSM), the latest digital mobile telephone technology, allows users to travel right across Europe and still stay in touch using a single handset and a single number.
- Cordless systems deliver high-quality speech and, because their transmissions are less powerful than GSM, the handsets can have exceptionally long battery life.
- Personal numbering services can provide nationwide mobility at a comparatively low cost.

BT is now taking the logical next step and creating a range of products that provide the best of all worlds. This article describes these *local area mobility* products and shows how they can lead the way towards the high-performance mobile services that people will demand in the next millennium.

The Cordless Market Today

Analogue cordless telephones

Since the first cordless telephones appeared at the beginning of the 1980s, they have become an accepted part of life in many homes. The telephones communicate with their base stations using low-power radio transmissions which restrict the distance over which they will work but increase the number of telephones that can be used in a given area. The first generation of cordless telephones, and the vast majority of those still in use today in the UK, are based on analogue technologies in which the speech is carried on radio signals between the handset and base station using frequency modulation (FM). Although improvements have been made to the analogue products over the years, they are susceptible to noise and interference due to the technology employed and the use of only eight radio channels. In some respects they have become victims of their own success—the chance of encountering interference has increased as the density of telephones in use has increased.

Digital cordless telephones

To address these problems, and to increase the number of applications for cordless technology, various second-generation standards based on digital technology have been developed around the world. In Japan and in the USA, these are the

Personal Handy System (PHS) and Personal Access Communications System (PACS) respectively. In Europe, the two digital standards implemented during the late-1980s and early-1990s are Cordless Telephony No. 2 (CT2) and digital enhanced cordless telecommunications (DECT). While these are all different technologies they use similar techniques to provide improved speech quality, security, capacity and functionality. They are intended not just as higher-quality versions of the analogue telephones, but also as technologies which can be used to implement more advanced ideas such as telepoint and wireless PBXs.

Telepoint

The telepoint concept is of a handset which can be used in the home, in conjunction with a base station connected to the telephone line, and which can also be used to make calls in public areas such as shopping centres and railway stations. In the UK, CT2 licences were issued to four operators; by the end of 1993 the last and most well-known of them, Rabbit, followed the others by closing down. The low-power transmissions restricted the coverage and initially the handsets could only be used on a single operator's network. The inability to receive calls in public places also contributed to an unflattering comparison with cellular networks. Elsewhere in the world, where the cellular networks have been less competitive, telepoint-like services have succeeded, most notably in Japan where over 7 million PHS handsets have been sold.

Business cordless

Until the advent of the second-generation digital standards, the take-up of cordless telephony in the business market was very poor because of the interference, noise and poor security encountered with analogue products. The situation has changed with the arrival of digital cordless which offers high-quality




Figure 1—Wireless PBX

encrypted speech and the full set of features normally provided by a PBX. In a typical business system (see Figure 1), a number of base stations are used to provide the area of coverage required. These are connected back to a cordless controller which tracks the location of handsets ensuring that calls are delivered to the right handset. The system passes calls in progress imperceptibly from base station to base station as the user moves around. Cordless controllers are available today which connect into existing PBXs by analogue or digital extensions. As far as the PBX is concerned the handsets appear as normal extensions. The controller and base stations effectively replace the wiring that would be needed to every extension in a fixed system.

DECT

Although DECT is the more recent of the two digital standards in Europe, it is expected to be the dominant technology for both the business and residential markets. It has the advantages over CT2 of a larger allocation of spectrum and a more advanced channel-allocation technique which allows more calls and more users to be supported in a given area. More importantly, volume sales within Europe have already reached levels where the handset prices are significantly lower than those of CT2, and this downward trend in prices looks set to continue. DECT uses encryption techniques to provide privacy over the radio link, and the handset continually monitors transmissions so that it can transfer the call to a better channel if interference is detected. As a standard it has great flexibility in that it supports




Figure 2—A single-line DECT cordless phone for the home

products ranging from single-line units for home use, through to ones for hundreds of users in a large workplace.

Although the products on sale today are primarily for speech services, DECT is capable of supporting a wide range of services with standards already defined for a variety of data applications. The standards are supplemented by a number of profiles designed to allow handsets and terminal equipment from one manufacturer to work with base stations and controllers from another. The profiles define restricted sets of functions that should be included for basic voice and data services. The generic access profile (GAP) defines the set of functions for basic speech services and will become mandatory for all new products introduced into the European Union from 1998 onwards.

Local Area Mobility for Business Customers

Business PBX systems employing cordless technology have been available in the market place for many years, typically providing mobility within one site, such as an office block or factory building. These have proved cost-effective where there is a defined requirement for single-site mobility. One such system was used in 1995 in a trial within BT,

Figure 3—Single site COSMiC

which involved 55 office workers at a multistorey block in Colchester. There were two novel aspects to the trial compared to systems on the market at the time: firstly, service was provided through external base stations mounted on poles (rather than base stations on walls in offices), and secondly, all control equipment was sited at a central point rather than on the office premises.

The explosion in use of mobile telephones has been one factor that has fuelled a demand for increased functionality from business cordless systems, to encompass multi-site roaming and automatic diversion to a mobile network when off-site. Roaming is where a user can visit other sites and receive incoming calls via their home site extension number on their cordless handset. The initial Colchester trial did not include any roaming or off-site mobility capability, and after exploring possible solutions with manufacturers it was decided to trial a new product known as the *mobility server* as part of the BT OnePhone development. The system came to be known as *COSMiC* (Cordless On-Site Mobility in Corporates) and offered an integrated cordless system and personal numbering platform, supporting dual-mode telephones, that operated as an extension on the Colchester cordless system when around the office buildings and a true mobile telephone elsewhere (Figure 3).

Sixty prototype dual-mode handsets were purchased for the trial, and the personal numbering platform ensured incoming calls would reach users regardless of which part of the handset they were using. All users had their directory number re-routed to the mobility server, from where the call would typically be sent to the cordless or mobile part of their handset, and on no reply to a colleague or to voice mail (see Figure 4).

Initially, users were offered access to the full facilities of the personal





Figure 4—Possible call answering positions and typical incoming call flow

number platform, such as multiple call-forwarding destinations through administration controlled profiles, call screening, and ubiquitous access from work, home or public network to swap

between profiles and activate facilities. At the same time, users were given the mobile extension (Mob X) facility on their mobile telephones, so that they could continue to dial other

The prototype dual-mode handsets were essentially cordless telephone circuitry and an interworking board sandwiched between the keypad and radio parts of a mobile telephone. When the handset is in range of a cordless system to which it has been pre-registered then it operates in cordless mode, but otherwise it operates as a conventional digital mobile telephone. Feedback has shown that these telephones are considered too bulky, leading to use of the smallest battery available and consequently poor battery life. A more integrated production model, which uses a single antenna and is only 1 mm larger than the standard mobile telephone on which it is based, is due soon.



Figure 5—Dual-mode handset

extensions with four digits regardless of the mode of their handset. This was tailored to mirror the Colchester numbering plan. Finally, users had a new voice-mail system to cope with, where many users had no voice mail at all beforehand. It was realised early on that introducing so many changes to service and so many new options at one go was confusing for the users. As the trial deliberately targeted heavy telephone users who were often out of the office for part of the week, it was rarely possible to get more than one or two people together at one time to explain the service. Although all the necessary manuals and quick reference guides were supplied there was a high overhead on the helpdesk staff in demonstrating and explaining the equipment face-to-face to users, determining and programming their optimum user profile, obtaining, delivering and activating their Mob X SIM card, and answering early user queries. This is an important lesson for future roll-out of such complex systems.

Colchester COSMiC was essentially a technical trial. When this concluded, further trials were instigated to test customer reaction to a tailored service package. Four

Showcase sites were set up across BT, where BT employees used the telephones on a daily basis and external customers could be brought to see the capabilities of the system. COSMiC at the BT Workstyle 2000 buildings at Stockley Park and Westside went live in the summer of 1997. A further two sites at BT Centre in London and BT Mobile's premises at Arlington Business Centre in Leeds went live during October 1997 as Phase 2 of the Showcase. BT Mobile operates as a stand-alone system, but the other sites are interconnected through 30-channel private circuits to allow cordless roaming, a key feature of BT OnePhone (see Figure 6).

Handsets from one of these three sites will register automatically as visitor handsets on arrival at one of the other sites. Incoming calls received at the home mobility server are forwarded to the visited mobility server for delivery to the handset. Outgoing calls are made through the home mobility server over the private circuit, which ensures that the user's home line is 'busy' when a call is being made so that the appropriate response can be made to subsequent callers.


Customer feedback from the Colchester and Showcase trials has been enthusiastic in many areas and has shown that the service fits well with certain work styles. A few users like to control the destinations for their incoming calls themselves and want access to all the facilities of the system, whereas most people are content with a pre-defined search path that requires no maintenance. The automated dual-mode handset switching between cordless and mobile is viewed as key to providing continuity of service between the fixed and mobile networks, and many users have taken advantage of the multi-site roaming facility.

The installations at the Showcase locations have highlighted a number of practical problems that need to be overcome for the final product. In particular, the present roaming architecture requires dedicated private circuits between systems which represents a considerable extra cost to the customer. Future implementations of the mobility server will address this issue, and there is an ongoing dialogue with the manufacturers to ensure that the best possible solutions are adopted.

Local Area Mobility for Residential Customers

The residential market has already embraced the concept of local area mobility, as shown by the volume sales of analogue, and now digital, cordless telephones. The current products generally consist of a single base station with an integrated handset charger connecting to a single line. Over the coming year, new DECT products will be introduced into the market place which will support more than one base station to increase the coverage area. Other enhancements will see units that connect to ISDN2 or multiple analogue lines. These products will still only provide the limited mobility of a cordless 'tail' from a fixed line connection in the home. The imminent launch of DECT/GSM dual-

Figure 6—Call and signalling flow for incoming call to handset visiting another site



mode handsets will break through this restriction to offer the prospect of a true nationwide mobility service on a single number while retaining the call quality and fixed line tariffs of digital cordless when at home. The single number element of this service will be enabled by automatic routing within the fixed network that delivers the call to the appropriate part of the handset over the appropriate network. Ultimately this routing function will be provided by the intelligent network for which the necessary developments are being standardised under the banner of *cordless terminal mobility*. In the meantime, existing IN functionality could be used to offer an earlier service, such as that used to provide BT's personal number service, OneNumber. However, in the first and simplest implementations, the base station will use Select Services at the local exchange to divert calls to the GSM number when the DECT handset leaves the coverage area. At the same time the dual-mode handset itself will notice that it has left coverage and will switch into GSM mode. ISDN2 base stations will be able to support a number of handsets, each of which will have an individual telephone number.

The Future of Local Area Mobility

All product launches contain an element of compromise. While everyone wants to see products with as much functionality as possible, there is also the need to get products out quickly to start to satisfy customer demand. OnePhone is no exception to this, and a number of ways in which the service could be improved have been identified.

Initially, as earlier sections have explained, the OnePhone business and residential products will be separate. From a customer's viewpoint, however, it would be better if there was no such distinction. In an ideal world, customers should be able to take their business OnePhone

handsets home and have their calls automatically diverted to their home base station (assuming, of course, that this is what they want). To satisfy this need, later releases of OnePhone will use network intelligence to bridge the gap between business and residential solutions, enabling BT to sell an integrated family of compatible OnePhone products.

There is no reason why the evolution of OnePhone need stop here. The use of digital cordless technologies provides opportunities to exploit another significant trend in the telecommunications market today—the rapidly increasing demand for data services. DECT, the digital cordless technology at the heart of OnePhone, is inherently capable of supporting data. In fact, the first DECT data product to be approved was a wireless local area network, and the standards now define support of data at rates up to 550 kbit/s. If this sort of capability is combined with the enhanced GSM data services that will appear over the next year or so, it will be possible to create a OnePhone product that will deliver bit rates in excess of 50 kbit/s when the user is out and about and in excess of 500 kbit/s when the user is in the office.

If even higher bit rates are required, alternative wireless technologies could be used. Today, wireless local area network products, capable of delivering data rates in excess of 1 Mbit/s, are readily available. Wireless infra-red links, capable of supporting 4 Mbit/s over a range of about one metre, are now included as standard in the vast majority of portable computers now being made.

It is also clear that the push towards higher wireless data rates is set to continue. HIPERLAN, the High Performance Radio Local Area Network standard, has just been ratified. This defines a radio system capable of delivering video, voice and data at 10 Mbit/s, over distances of around 50 m. So, by combining

HIPERLAN and enhanced GSM, a very powerful next-generation OnePhone product, capable of high data rate multimedia support, could be created. As well as being an attractive proposition in its own right, such a product would also start to deliver some of the elements of the universal mobile telecommunications system (UMTS)—the European vision of communications in the next millennium.

Conclusions

Many different telecommunications products are now trying to satisfy the high demand for mobility, including:

- mobile telephones (such as GSM),
- personal numbering, and
- cordless systems.

The OnePhone concept takes elements from all of these products and builds upon them to form integrated offerings that deliver 'the best of all worlds'. After extensive analysis and trials, the first OnePhone product to be launched will be aimed at the corporate market. This will allow business users to exploit low-cost high-quality cordless communications when they are in the office, and the ubiquitous coverage of GSM when they are elsewhere, while being contactable via a single fixed network number.

This initial OnePhone product will be followed by others aimed at the small business and residential markets. As the market potential for this product has been predicted to be as high as 4 million units in the UK alone, it is clear that this could be a true mass market product.

Subsequent OnePhone offerings will look to integrate business and residential offerings. By taking this approach, and by enhancing OnePhone's ability to support data services, OnePhone can become an important step towards the UMTS—Europe's vision of telecommunications in the next millennium.

Acknowledgements

The authors gratefully acknowledge the contribution of Ericsson and colleagues in BT Mobility Solutions towards this article and the various trials and products which it describes.

Biographies



Steve Buttery
BT Networks and
Systems

Steve Buttery joined BT in 1981 as an apprentice. He worked on installation for two years before leaving for university as an award holder. Upon graduating in 1989, he joined the Radio Division at BT Laboratories. Since then, he has worked on several radio-related projects, ranging from writing software planning tools to supporting the development and testing of a rural radio access system. Currently, he is the Resource Manager for the Spectrum Management and Radio Team.



Bob Merrett
BT Networks and
Systems

Bob Merrett joined BT as a student apprentice in 1965. After graduating from the University of Wales with a B.Sc. in Electrical and Electronic Engineering in 1969 and a Ph.D. from the University of Cambridge in 1972, he joined the reliability division at BT Laboratories where he led teams assessing the reliability of semiconductor packaging, reed relays and a range of discrete electronic devices. In 1983, he joined the Devices Division and led a team designing high-speed GaAs integrated circuits and then a team developing InP OEICs. Since 1990, he has worked on various aspects of radio local loop and currently coordinates the technical projects on wireless access and local area mobility.



Adrian Pote
BT Networks and
Systems

Adrian Pote joined BT in 1987 after graduating from Oxford University with a B.A. (Hons.) in Physics. He has since worked on a variety of radio-related projects including microwave component design, radio-over-fibre transmission and access radio. He began working on cordless systems in 1995 and now leads the OnePhone project in the Spectrum Management and Radio Team. In 1994, he received an M.Sc. in Telecommunications from Essex University and is a Chartered Engineer and a member of the IEE.



Andy Sago
BT Networks and
Systems

Andy Sago joined the Post Office Research Department at Martlesham Heath in 1975 as a Trainee Technician (apprentice). At the end of training he chose to settle in the Radio Division before gaining a Minor Award to study at the University of Essex for an Honours Degree in Electronic Engineering. Since returning, he has worked at BT Laboratories on trunk radio transmission systems, wireless local area networks and radio local loop technology. After undertaking the technical trial for the initial dual-mode handsets and installing COSMiC systems within BT, he is currently involved in bringing the COSMiC product to market as the first part of the BT OnePhone family.

The Underground and Copper Access Network Infrastructure

BT's underground infrastructure and associated copper network is the fundamental building block which supports service delivery to the customer. This article briefly outlines the issues associated with ensuring the individual components are fully capable of meeting current and future demands.

Introduction

The BT underground and copper access network is a key asset to the company. In approximate terms it comprises:

- 0.5 million kilometres of duct.
- 4.5 million underground structures.
- 3.5 million overhead poles.
- 2.0 million kilometres of cable.
- 87 thousand street cabinets

All of these serve some 30 million working lines. With a net present value of some £4.5 billion, the scale and size of the network are easily appreciated and its complexity readily imagined.

This article reviews the recent developments which have been carried out to ensure the extensive civil engineering infrastructure and the copper network remain fully operational and capable of delivering advanced services to customers for the foreseeable future.

Various network components are discussed individually with a view to providing a brief appreciation of the technology being applied to maintain and improve the performance of the network.

Underground Structures

Underground structures have been deployed by BT since the late

nineteenth century and have evolved from simple buried brick box units to complex multi-storey reinforced concrete structures. These structures, whether they be simple joint boxes or complex entities, are the physical entry points to the core and access network allowing staff direct intervention to the network for maintenance, repair and build operations.

The current versions of these structures have been structurally designed to meet stringent design requirements stipulated by European, British and Department of Transport standards such that collapse does not occur when the structure is exposed to the increased forces generated by modern traffic. To achieve this, the structures have been designed as ridged units which deploy the inherent stiffness of the roof and wall elements to support the loads. This methodology has led to cost-effective thin-walled structures of significant strength, easily able to accommodate present and anticipated future traffic forces. By fully utilising the permissible stress envelopes for the constituent materials, and arranging the high-stress elements such that forces are distributed throughout the structure, a very strong thin shell has effectively been created which is very durable and resilient.

The installation of these structures varies from the simple, for a standard joint box, to the very complex, for a large buried structure in a city centre. Because excavations are dangerous, great care is taken to



Figure 1—Typical city centre site showing traffic diversions, protection and fencing off



Figure 2—Temporary support of existing BT cables




Figure 3—Schematic of old brick structure.



Figure 4—Schematic of replacement concrete structure

ensure that the safety of staff and the public at large is taken seriously, with work sites being protected and, in the case of large sites, securely fenced off (Figure 1). Where existing cables are discovered these are suitably supported to ensure no service failures are introduced during construction operations (Figure 2). Again, great care is taken to avoid damage to the existing plant.

BT has an ongoing programme to replace old structures with current designs where it has been identified that the older structures require upgrade or repair. Generally this is done after inspections have identified that there is a need to enlarge, replace or relocate the structure. In the majority of cases a simple replacement of the existing structure with one from the current range of standard structures is adopted. However, in certain situations, where no standard structure is available, a specific special structure is designed by the Civil Engineering Policy Team (or its agents) for a particular situation. Typically, an old complex brick structure (Figure 3) needing replacement is carefully measured and a new concrete design

(Figure 4), together with construction details, is developed. These details are then sent to the field for the construction process to begin.

The team is currently developing proposals for a completely new infrastructure platform comprising prefabricated elements which will be delivered to site ready-made and can be installed easily. Substantial productivity gains are predicted and unit costs for the individual elements should be reduced.

While the above overview discusses the broad issues associated with underground structures, the detail and complexity of developing uniform standards and practices can be easily appreciated.

Ducts

Some 0.5 million kilometres of duct form an essential part of the network. The BT duct networks are the distribution arteries through which the majority of buried cables reach the customer.

Visually, ducts appear as simple plastic tubes, but this betrays the complex technical issues which are considered when producing speci-

cations to ensure they are capable of withstanding fully the severe conditions they are expected to survive during their 50-year design life. Rigorous stress and materials analysis is undertaken so that the product will withstand crushing forces, from initial backfill compaction to those generated by vehicular traffic.

Materials technology also plays a significant part in the development of duct products, ensuring they are resistant to expected chemical attack, service temperature variations and accidental physical abuse.

The BT duct infrastructure is of fundamental importance to the efficient management of the underground distribution network and a key asset to the business in its drive to support business growth. Once buried, ducts may be unseen but they are certainly not forgotten. Using the human body as an analogy, main arteries and veins are generally hidden but if blockages occur the whole body soon realises the potential for breakdown.

Closures

Closures are essential elements of the network in so far as they are the components which ensure made-up cable joints are protected from the environment and stop water penetration, which is a major fault liability in unprotected joints.

Work is continually ongoing to improve the performance of these closures and take advantage of the latest industry developments.

Recent introductions to the network have included the Sleeve 31J series. The need for this product was identified as a replacement for jointing posts which had been extensively used on numerous




Figure 5—Typical jointing post

housing estates during the 1960s and 1970s. The jointing posts (Figure 5) were not effectively sealed, offered poor security and access was often impeded by changes in ground level and cutting back cables. They had clearly outlived their useful working life and were rapidly becoming a serious fault liability.

Sleeve 31J (Figure 6) was introduced as an easily installed replacement which offered a secure sealed environment and could be installed without the need to break down the joint and remake it. Great care was taken to ensure installation practices were compatible with Sleeve 31 series parameters, so eliminating the need for training.

Another product introduced into the network is a new airblock kit, for use on the pressurised exchange side parts of the network, which offers significant productivity and installation advantages. Previously, installing airblocks was a cumbersome procedure entailing the use and disposal of resin packs, which required time and care during installation.

The new airblock kits (Figure 7) have the blocking compound, in the



Figure 6—Replacement Sleeve 31J

form of solid pellets, contained in a porous bag which is positioned around the cable pairs and liquefies when heat is applied to the heat-shrink sleeve. This product has led to a 50% productivity gain while eliminating the use of resin materials which have to be carefully managed.

Copper Connection

Copper connection systems are specified such that the integrity of any copper pairs within a cable system is not degraded and that maximum signal performance is maintained along the entire route of any individual pair.

The copper network is the major service-carrying component within the BT network and has proved a resilient asset to the company. Over the past decade significant improvements have been carried out to ensure that this valuable asset is capable of carrying advanced digital services and making it a viable digital platform for the future. The emphasis has been to future-proof the copper network and its compo-

nents to maximise the service delivery requirements that customers demand.

As part of this copper network enhancement, it was soon realised that, while the actual copper wires easily handled modern transmissions, the performance of existing connection systems would suffice in the short term but deteriorate rapidly due to the new demands being placed upon them.

Following a comprehensive review, the decision was taken to introduce a new range of connection systems for use in the network.

The first product to be introduced was the modular block terminal (Figure 8). Block terminals are mounted on poles and walls, and act as distribution points linking the primary connection points to dropwires serving individual customer premises. They are usually the last flexibility point on the network. The modular block terminal replaced the BT76 and BT86 series and incorporated a new insulation displacement connection

Figure 8—The new modular block terminal showing housing and insulation displacement connectors





Figure 7—New airblock kit showing blocking compound and heat-shrink sleeve



(IDC) system which would meet present and future performance requirements. These units have been introduced on a phased basis within new and replacement work programmes.

Essentially the new block terminals come as 20- and 10-pair capacity versions and the housing has been designed to allow venting and to provide enhanced security features. The actual IDC modules are single entities which are fitted into the housing on a pair-by-pair basis dependent on the number of drop wires. The IDC modules are easily installed and come with a reusable test and lightning protection facility, in addition to enhanced electrical performance.

The next step in enhancing the copper network is the use of modular jointing units for new and remade cable joints housed within the buried network. Modular jointing has been used on the pressurised exchange side (E side) of the network for some years but is currently being rolled out in the distribution side (D side).

Modular jointing (Figure 9) comes as a complete package which includes three basic strip connection modules being 10-pair straight, 10-pair half tap and 10-pair pluggable units which house the insulation displace-

ment connectors and are designed to meet all jointing requirements.

To allow the benefits of modular jointing to be fully exploited, a modular jointing tool kit is available which comprises a storage case, a splicing head, a hand presser tool, a single-wire insertion tool, a single-pair test probe, a joint support bar with clamp and a user manual. The introduction of modular jointing gives significant productivity gains and the roll out is supported by a comprehensive training programme for field staff.

Work is currently taking place to introduce high-density modular cross-connection strips for use within street cabinets. These strips will increase cabinet capacity by some 30 per cent and meet present and future service requirements.

The connection of copper wires is a complex issue but the systems currently being introduced will cater for the future demands being placed onto the network.

Conclusion

This brief overview of BT's underground copper network has highlighted its importance to the business and the care and attention it receives to ensure it remains fully utilised, maximising revenue from its usage. All its constitu-

ent components are carefully specified, evaluated and approved by a dedicated team of professionals who remain convinced that a well-managed copper platform has great potential for contributing to the ongoing success of the company.

The platform may be buried but it is most certainly alive and well.

Biography

Mazz Da Rios
BT Network and Systems

Mazz Da Rios is a Chartered Engineer and holds a Masters Degree in Business Administration. He joined BT in 1987 as the project manager in the Radio Works Group and was involved in the construction of the West of Scotland radio link entailing construction of five mountain-top sites. He moved to the External Engineering Division in 1989 where he has led teams involved with civil engineering, government regulatory issues on the New Roads and Streetwork Act, overhead-network infrastructure, copper-connection systems, buried duct and closure systems including street cabinets. He is currently the business coordination manager for BT's External Engineering division overseeing the commercial, project and delivery issues associated with the work of the unit.

Figure 9—The modular jointing system



Fighting the Information Terrorists

Information is one of an organisation's most precious assets. The explosion in electronically stored information and the extensive networking of computers has placed that information under new threats. This article examines some of these threats, trends in information engineering and reviews techniques for protecting information.

Introduction

Over 30 years ago the United States Airforce ROTC Manual, *Fundamentals of Aerospace Weapons Systems*, defined a military target as . . . 'any person, thing, idea, entity or location selected for destruction, inactivation, or rendering non-usable with weapons which will reduce or destroy the will or ability of the enemy to resist'. Thirty years ago it might have seemed somewhat perverse to have included an 'idea' as a legitimate target for the USAF. Today, however, we are entering an era of information warfare where hostilities are conducted in a non-physical world. The computing and networking paradigms have shifted from secure networks with unprotected data towards a greater use of insecure, open networks with protected data (see Figure 1). This increasing trend towards networked computers and platforms presents

unrivalled opportunities for the technically astute enemy.

Those enemies are also likely to come from diverse backgrounds. Today's systems are at risk from a spectrum of assailants ranging from both the casual intruder to the organised efforts of national intelligence agencies. We often have an image of the computer hacker as some socially maladjusted teenager clothed with acne and an anorak who spends endless hours trying to penetrate other people's computer networks. While such individuals undoubtedly exist, a far more serious threat is posed by those who would manipulate these casual intruders. Today's information terrorist recruits members from across the globe.

When information can be distributed on a global scale with a single keystroke we must also be mindful that this is true of misinformation. The recent announcement that a 40-bit secret key crypto-system had been broken by a student at Berkeley in just 3½ hours showed just how insecure 40 bit systems are. However, within a very short time certain people made statements trying to pour scorn on this achievement and making the preposterous assertion that 40 bit secret key systems were perfectly secure. Consider also the damage that could be caused to a company by a widely-publicised false assertion of insecurity. The losses consequent on such a damage to reputation are potentially enormous and difficult to quantify.

Information is something of a nebulous concept, yet it is one of the most valuable assets of any company or organisation. The protection of that information is a complex task. While we shall concentrate on the

Figure 1—The change in threat model as the computing and networking security paradigm has shifted. The rewards and damage caused by criminal abuse of privileged information are much greater than the potential damage caused by a breach of security at the physical layer




Figure 2—The function of a firewall.
Only authorised access is permitted


technical threats and defences it is important to realise that computer systems are operated by human beings. Some very sophisticated technical protection mechanisms have been neatly side-stepped by the simple expedient of obtaining passwords over the telephone under false pretences.

Understanding an organisation's security requirements and implementing a sensible security policy is a difficult task. Often security specialists talk in terms of a layered approach to security in describing security requirements. Each of these security layers is taken very seriously by BT, and teams of specialists operate to ensure that security at every layer is as good as it possibly can be. This article considers a simplified security model in which there exist only four layers as displayed in Figure 1. As we shall discuss, compartmentalisation of an overall security framework in this fashion can be artificial. It does, however, provide a convenient classification of possible threats and countermeasures.

Physical Security

The first layer of an overall security framework is traditionally the physical. That is, concerned with access to equipment and/or safe sites. Physical security can be achieved through a variety of mechanisms for access control ranging from simple pass-card procedures to advanced biometric techniques, or indeed any combination of these. One of the most promising new techniques to have been developed uses iris-scanning methods to identify individuals. Other techniques combine the use of smart-card technology with finger-print identification so that the smart card is activated by a positive biometric verification.

To some extent it is artificial to separate out the security layers as we have done in our simple model. For example, if we consider the case where access control to a secure site is



achieved by interrogation of a database, then it is crucial that the server containing the database records can not be compromised or its information obtained by another route. It is also important that server failure, for whatever reason, does not compromise the security of the system.

Businesses of today are becoming more reliant on advanced high-speed information networks, and controlling access to those data networks is a primary concern. Physical security, while still an important part of the overall security framework is, perhaps, becoming less of a challenge as more and more of an organisation's resources and capabilities are migrating onto distributed computer platforms. Control of access to those resources is an important function of security at the network layer and is achieved through the use of firewalls.

Network Security (Firewalls)

The first point of defence for any organisation's information networks must be at the points where those networks connect with an external network or system. If an unauthorised intruder cannot access those internal networks, then the only threat to that information is from within. Preventing unauthorised access by an external source while allowing authorised access to external systems for internal users is a non-trivial technical challenge. BT has put in place sophisticated technologies for meeting that challenge and has a comprehensive range of network tools to control access to and from external networks. The most important technologies for

achieving this goal are known collectively as *firewalls*. The job of a firewall is to implement an organisation's security policy with respect to connections between internal and external networks. In a sense, firewall technologies span the network level of Figure 2 with some overlap into the physical and system layers.

The simplest, and safest, security policy for an organisation can be stated as '*deny everything which is not expressly permitted*'. This kind of policy should be implemented wherever possible. However, defining what should be permitted is not an altogether trivial task and regular review of policy requirements is essential.

A firewall usually consists of dedicated hardware or software running on a router. The location of the firewall is, obviously, at the interface between internal, trusted networks and external, untrusted systems. There are several technical components that are used for building firewalls:

- *Packet filtering* gives the ability to analyse the incoming and outgoing packets and allow or deny access.
- *Application proxies* work at the level of the application protocol rather than the network and transport layer. For example, a World Wide Web (WWW) proxy is able to decode the contents of a WWW packet so that it can be forwarded by the proxy to its intended destination. These proxies are processes run on a machine, known as a *bastion host*,

Figure 3—Typical proxy server firewall configuration showing the demilitarised zone. A proxy server is able to tightly control connections to external services

that is situated in an area of the network known as a *demilitarised zone* (see Figure 3).

- Address or port translation allows the details of the internal network to be hidden from outside of that network by using an address translating gateway. This function can also be achieved by proxies, but use of an address translating gateway can overcome some of the limitations of using proxy servers¹.


These techniques can be combined in various ways to give protection to an organisation's internal networks. More details of firewall technologies can be found in Reference 1.

A firewall must be the first level of defence against hostile intrusions of an internal network. However, what happens when an internal network is penetrated or when sensitive information must be transmitted across a hostile network environment? To protect against these threats encryption techniques must be used along with suitable access control mechanisms for internal databases.

System Security

At the next level of security individual systems must be safeguarded. Protection mechanisms include simple password verification for access to individual PCs or servers. Security policies at this layer are crucial, as is the effective training of personnel. Proliferation of password protected systems each requiring separate login identities and verifications is often counter-productive, leading to a situation where individuals write down a growing, and frequently changing, list of passwords and store this list in an easily accessible place. BT's security policies in this regard are carefully thought out to provide the maximum security combined with ease of use.

Once again biometric techniques can be exploited to provide a greater level of protection combined with a greater ease of use. Analysis of typing



patterns at a keyboard, for example, can provide a unique individual characteristic. Combined with password protection this can be a simple and reasonably effective tool for protecting access to a particular PC or network. Other biometric techniques such as iris scanning and fingerprint identification could also prove useful, as could smart-card technologies.

Care must also be exercised on the use of individual PCs connected to a network. It has been known for individual machines when, quite innocently, running certain software applications on a network to breach that network and 'punch a hole' in the firewall protecting the network. Fortunately, these breaches are rarely serious and are normally captured by the firewall's monitoring tools. However, it serves to emphasise once again how security at one level can have an impact at another security layer:

Information Security (Encryption)

We have so far concentrated on the protection of information as a vital corporate asset within an organisation's internal networks. As we move towards a global information infrastructure (GII) information will increasingly be seen as a rich target, whether that information is contained within an organisation's own networks or whether it travels on open, public networks. The threat model is moving further along towards the information layer of Figure 1. While the rich connectivity of today's computer networks will bring about unprecedented benefits it is that very connectivity that increases the level of threat. There are a number of potential threats to information on open networks²:

- interception and consequent violation of secrecy or privacy,
- impersonation to shift liability,
- impersonation to originate fraudulent information,
- modification of legitimate information,
- fraudulent authorisation of transactions,
- denial of responsibility,
- false claims to have received information,
- false claims to have sent information,
- denial of receipt or time of receipt,
- enlargement of legitimate licence,
- modifying the licence of others,
- concealing the presence of information in some other message,
- denial of time critical information,
- presence of undetected relay points,
- traffic analysis,
- revelation of secret information to undermine integrity protocol,
- corrupting the function of software,
- violation of protocol by introduction of false information, and
- undermining of confidence by causing protocol failures.

Advanced cryptographic techniques, properly implemented, are necessary to effectively combat these threats³. BT has considerable experience of using cryptographic techniques to protect the data on its internal networks. However, these cryptographic techniques can also restore the confidence in the use of open networks. The potential benefits of a GII are immense as are, correspondingly, the opportunities for abuse. That abuse can be minimised and made negligible by the proper use of cryptography.

There are two very important, but distinct, services that can be provided by the use of cryptographic technologies. These are confidentiality and authentication. The most widely-known function of cryptography is that of confidentiality or secrecy. In this case a message is scrambled to make it unintelligible to any but authorised recipients. However, perhaps the most important application of cryptography is to provide authentication. In many cases it is more important to know where and when a document has come from and that it hasn't been changed than to simply keep it secret.

In traditional cryptographic systems two communicating parties must share some secret in order to be able to send subsequent messages with any degree of confidentiality. This secret, usually a sequence of binary digits, is known as a *key*. The key is used to 'activate' a mathematical transformation, or algorithm, which scrambles the message so that anyone intercepting the scrambled message cannot recover the original without the key. The obvious analogy is with a lockable box (Figure 4). The message is placed in the box which is then locked using the key. The box can then be sent in the knowledge that only those recipients with the key can unlock it and recover the message. The algorithms (that is, the locking mechanisms) if well designed can be publicly known without compromising the security. In fact it

is a recognised design criterion for building a cryptographic system that the security should not reside in the secrecy of the algorithm but only in the secrecy of the key.


There is an obvious flaw with these kind of systems: all of the legitimate parties in the communication must share the same secret key in order to be able to read the messages. To overcome this problem a revolutionary new form of cryptography system, known as public-key cryptography was proposed in the mid 1970s. As we shall see, systems of this kind turned out have much greater implications for communications than just the provision of secrecy. In public-key cryptography each person has two keys. One of these keys is publicly available, the other private. In terms of the lockable box analogy the public key, which everyone has access to, is used to lock the box which can then only be opened by the person with the corresponding private key (see Figure 4).

With public-key techniques two users wishing to communicate securely do not have to have exchanged an initial secret. However, it is impossible for two people who have

never met in any fashion to communicate with any security. The simple reason being that a third person could be intercepting all messages and impersonating the legitimate parties. In traditional cryptography, the secret key exchange is the 'meeting' between the parties. In public-key systems, the publicly available information, the public key, has to be trusted or to come from a trusted source. The common point of trust acts as a 'meeting' point between the two parties for public-key systems.

The real power of public-key techniques, however, is not simply in providing confidentiality. In just the same way that we wouldn't accept a cheque without an appropriate signature, we shouldn't accept any electronic data without some form of authentication. Public-key techniques allow us to 'sign' data electronically using the concept of a digital signature. A digital signature binds the identity of the signer with the document that is signed. Once data has been signed, it cannot, therefore, be altered without the signature being invalidated. The data itself can be trusted and can be stored or transmitted even over open, insecure networks.

Figure 4—Alice sends a confidential message to Bob using public-key techniques. Firstly she retrieves Bob's public key from a trusted public source. The message to be protected is then locked in the case using Bob's public key. Once the case is locked only the private key can open it so that only Bob can read the message intended for him.



The process of creating a digital signature is depicted in Figure 5(a). The verification process is depicted in Figure 5(b).

In order to be able to verify a signature, the signer's public key is used to unlock the encrypted hash. This procedure requires that the signer's public key must be trusted. To establish that trust, a user can get a public body, known as a trusted third party (TTP), to sign his/her keys digitally and issue electronic certificates verifying those keys. In order to verify the signature, the public key of the TTP must be used to check the certificate. Of course the public key of the TTP may itself be signed by a higher authority. This leads to the concept of a hierarchy of TTPs who provide certification services. Provided this trusted infrastructure is in place, digital signatures provide a much stronger mechanism than their handwritten counterparts. As our computing paradigm moves towards insecure, open networks and away from private secure networks, we shall find that digital signature techniques become indispensable.

With an infrastructure of TTPs and the use of digital signature techniques, we can now envisage transferring all of our paper-based processes into an electronic environment. The use of BT's intranet has been growing at an exponential rate and the cost savings made by using more and more electronic services have been enormous. Organisations and individuals can expect similar cost and efficiency savings with the introduction of advanced services over a trusted GII.

Future Issues

The ability to transfer paper-based processes into an electronic environment will open many new markets and allow companies to offer services that were previously unthinkable and possibly outside of their traditional business focus. Furthermore, these new marketplaces will be global. Such globalisation will undoubtedly cause many contractual and legal difficulties.




Figure 5(a) Digital signature generation. To sign a message digitally, a copy of the message is taken and a digest, or hash, of the message is produced. This hash is then encrypted with the signer's private key. The resulting encrypted hash, which is the digital signature, is appended to the original document.




Figure 5(b) – Digital signature verification. The encrypted hash is decrypted using the signer's public key. This hash is then compared with a hash calculated from the document. If there is agreement the signature is accepted. Even a slight alteration to the document is sufficient to cause the two hash values to differ significantly.

For example, providing an on-line service to offer financial advice to clients may have to operate under different regulatory regimes in different jurisdictions. When there is a dispute it may not immediately be obvious which jurisdiction takes precedence.

A related question, and one which is central to the whole TTP debate, is the issue of liability. The word 'trust' has been much abused and is usually wrapped up with personal notions of trust. For example the questions 'do you trust your bank?' and 'do you trust your best friend?' are using 'trust' in two very different contexts. We trust banks insofar as they are liable for errors. If an organisation has a strictly limited liability then we can not assume a trust relationship over and above this liability limit. In the TTP debate in the UK the question 'would you trust your bank?' is often used in a pejorative and wholly inappropriate sense. Clearly we only trust banks in respect of certain

transactions and then only insofar as their liability extends.

If hierarchies of trust are to become the new electronic business paradigm users will need mobile trusted anchor points (TAPs) into the system. Such devices are conveniently provided by smart cards. Trusted data can be loaded onto a multi-application smart card together with the appropriate public key certificates. TAPs are required, especially where mobility and the dictum 'anytime, anyplace, anywhere' are important drivers. When using a public electronic kiosk or another networked computer, which cannot themselves be viewed as TAPs, it is important to have some confidence in the transactions that are being conducted. Smart cards can act as TAPs and restore that confidence in the system.

If customers are to have confidence in any global system there must exist an appropriate confidentiality framework. That is, users must have the confidence that their transactions, when

so desired, are conducted in confidence and cannot be read by unauthorised persons. However, cryptography for confidentiality is viewed as a 'dual-use' good and is subject to stringent export regulation by most countries. The result of this has been to limit the strength of commercially available cryptography for general use. Unfortunately this limit is no longer adequate for material of even the slightest sensitivity. It is a moderately trivial task now for an organisation or individual with access to modest resources to break exportable cryptography on timescales ranging from a few minutes to a day. The security facilities built into exportable versions of software are therefore inadequate and should not be used to protect sensitive information unless there is absolutely no other alternative.

There are obvious national security issues to be addressed over the use of strong cryptography, as there are issues concerning law enforcement agencies' abilities to conduct surveillance operations. However, businesses and individuals have a legitimate and growing requirement for access to strong cryptography. In order to balance these conflicting requirements governments have proposed key escrow systems in which cryptographic keys for confidentiality are held by a third party. The third party would release these keys to the appropriate authorities upon provision of a legal warrant. It is now clear that some, if not most, of the export controls on strong cryptography will be relaxed if suitable key escrow mechanisms are in place. It is fair to say, however, that these key escrow mechanisms have not been universally welcomed.

Clearly the success of any future global framework for electronic commerce depends on the interoperability of the systems used. This, at the very least, involves agreement on worldwide standards. Certain standards are beginning to emerge and are likely to be widely adopted. For example, the X.509 certificate standard is the prime candidate for use in signature and public key applications. A crucial

advantage of the X.509 standard is its dependency on a hierarchical framework with root certification authorities. This centralisation of any system is important if customers are to have confidence in any certification process. Any certification of value bearing data must be appropriately underwritten and liabilities clearly understood. A certification structure containing a relatively few licensed certification authorities would provide these necessary customer assurances and maintain an underwritten service with transparent liabilities and a well-defined audit trail.

Any company must protect its assets if it is to succeed in today's highly competitive marketplace. Nowhere is this more evident than in the exciting world of telecommunications where developments occur on such a rapid time frame. Difficult though it is to quantify, information must be properly thought of as an asset and protected accordingly. Globalisation of telecommunications networks and the emerging global electronic marketplace force new demands on how we protect that information. New cryptographic techniques and protocols must be implemented to achieve the requisite level of security. BT is at the forefront of moves to ensure the security of this new electronic marketplace and establish new exciting electronic services for its customers.

References


1. HUBBARD, S. D.; and SAGER, J. C. Firewalling the 'Net'. *BT Technol. J.*, April 1997, **15**(2), p. 94.
2. SIMMONS, G. J. Contemporary Cryptology: a Foreword. In *Contemporary Cryptology: The Science of Information Technology*, ed. SIMMONS, G. J., IEEE Press, New York, 1992, p. vii.
3. PHOENIX, S. J. D. Cryptography, Trusted Third Parties and Escrow, *BT Technol. J.*, April 1997, **15**(2), p. 45.

Biographies



Simon Phoenix
BT Networks and
Systems

Simon Phoenix graduated from York University in 1986. He joined BT in 1989 after having completed his Ph.D. in theoretical quantum optics at Imperial College London. His research work at BT has involved fundamental investigations into the quantum nature of light and its application to telecommunications. After a period developing several new quantum cryptographic techniques and concepts, he moved into the security futures team at BT Laboratories where his interests now include commercial applications of cryptographic techniques and trusted third party services.



John Regnault
BT Networks and
Systems

John Regnault joined BT in 1971 and throughout his career has specialised in telecommunications research and development. Over this time he has worked on a range of communications technologies and has published 35 technical papers. He currently heads BT's security research group at BT Laboratories. His team of 16 security specialists are responsible for advising others across the BT group on new security threats and opportunities offered by advances in technology. Outside BT, John is an active radio amateur specialising in home construction and moonbounce communications at VHF frequencies. He also writes technical articles on new and advanced technology for radio magazines.

Beyond the Handset: Audio for the Next Millennium

Developments in communications technology, such as audio-visual conferencing and telepresence, have created new demands for audio delivery systems. In the past, a telephone handset, with its limited transducers, was adequate for simple spoken communication; however, future multi-user communication systems and immersive machine interfaces will be greatly enhanced by using high-quality three-dimensional audio. Such three-dimensional sound systems can position individual talkers' voices in space around the user, creating a greater feeling of immersion and improving intelligibility. This article discusses spatial audio technology and some advanced applications.

Introduction

Since the invention of the telephone in 1876, we have become used to communicating with people remotely by holding a handset to one ear. This is perhaps a rather unnatural method of communication. In a typical person-to-person conversation, we see the other person's facial expressions, pick up fine intonations in his or her voice, read his or her body language, use visual aids and hear a wide variety of sounds around us. In a situation where there are many people around us (such as in a cocktail party or business meeting), our hearing draws attention to those not in our field of vision who are conversing with us. We then turn to them to continue the conversation. The cocktail party effect describes our ability to tune into a conversation of interest in a room full of people talking.

The spatial sound cues we subconsciously interpret help us to interact with the world around us. For example, if you were standing in the road and you heard a car coming from behind, you would quickly get out of the way. The same cues are an essential part of our natural inter-person communication process alerting us to activities nearby and helping us with orientation and navigation.

An ultimate aim for a communication medium is to provide a mechanism for people to interact naturally to the extent that they are so involved in their task that they are unaware of the physical system altogether.

A very important component of a natural and compelling communication medium is realistic sounding audio; that is, audio that envelops a

person. This article considers a selection of methods for delivering realistic three-dimensional audio (some of which have been developed by the authors specifically for communications applications) and describes some example systems developed at BT Laboratories.

Spatial Audio

It is important to clarify the distinction between true three-dimensional spatial audio and surround sound as often found in cinemas and domestic audio-visual equipment. With spatial audio, sound can appear to originate from anywhere in the room and can move around the room irrespective of the loudspeaker's location. Generally, surround-sound recordings use the rear channels for effects (often with the same, or very nearly the same, signal coming from both channels), and speech is mainly reproduced from the front-centre channel. Unlike spatial techniques, sounds tend to originate from the individual loudspeaker locations.

Transaural

Imagine a system with only two loudspeakers placed in front of you and yet you can hear sounds coming from different positions around you, even from behind. Does this seem impossible? By using some special processing of the audio, a true three-dimensional sound field can be created from only two loudspeakers. The system, known as a *transaural* system, works because of the way humans determine the location of sound sources. Firstly, sounds take a different amount of time to get from the original source to each of our two ears. The delay introduced between




Figure 1—Transaural processing with cross-talk cancellation

arrivals is known as the *interaural time delay* (ITD). The ear nearest the source receives the signal first; the signal is also louder at the ear nearest the source, known as *interaural intensity difference* (IID). Secondly, and more complexly, sounds arriving at our ears from different directions bounce off different parts of the outer ear and are passed into the ear canal. The irregular surface of the ear causes slight changes to the sound from which we are able subconsciously to interpret directional information. By making measurements, it is possible to understand what the ear does to a given sound from a particular location. This information can then be used to modify a sound to trick the ear/brain into thinking that the sound comes from a given direction. The way the outer ear modifies a sound for a given relative source position is known as the *head-related transfer function* (HRTF) and for a fully three-dimensional system a wide range of HRTFs needs to be determined.

Transaural systems reproduce the desired sound direction by taking account of the fact that each ear will hear both of the loudspeakers. Crosstalk cancellation is used to ensure that each ear receives only the desired sound. It works by changing the signal so that the sound arriving at the left ear from the right loudspeaker is cancelled out, but the sound coming from the left loudspeaker is left intact. The same effect occurs at the right ear to ensure that it only hears the right loudspeaker and not the left one. A block diagram showing crosstalk cancellation in action is shown in Figure 1.



Figure 2—SmartSpace chair using transaural processing

Transaural systems work well, although it is important that the listener is precisely located at the correct position relative to the loudspeakers. Transaural processing has successfully been used in the BT SmartSpace chair where the user's head position is known in advance and does not move very much while undertaking a communication task.

The SmartSpace chair (as shown in Figure 2) combines a video screen, computer terminal, video camera and spatialised audio. The chair has been developed by BT to suggest an alternative to the traditional office desk. The video screen wraps around the user to provide a feeling of visual immersion for use in telepresence applications. The chair currently has two loudspeakers mounted below the screen at either side. By using transaural processing, a 'sound bubble', giving the feeling of complete audio immersion, can be created around the user.

Binaural

Binaural systems are simpler to generate than transaural because the user wears headphones, eliminating the need for crosstalk cancellation. Users do not generally like wearing headphones while they are working, and headphones tend to detract from feeling immersed in a natural communication experience. However, there are situations where the use of headphones may be considered desirable or essential; for example, someone engaged in a meeting while travelling by public transport.




Figure 3—Binaural signal with and without head tracking

Headtracked binaural

Because binaural signals are heard with headphones, spatial sounds will move when the head moves. In communication systems it may be desirable to keep the sound location static even if the listener rotates his or her head. By tracking the head movement and applying different HRTFs for different head positions the apparent sound source location can be kept static. A typical practical implementation uses 128 different HRTFs, corresponding to each 2.8 degrees of head rotation.

The modification of the HRTF can be performed in real time using digital signal-processing equipment. Figure 3 shows the effect of head movement both with and without head tracking applied. The benefit of head tracking is that audio image locations remain static during movement of the listener's head; the disadvantage is the requirement for real-time digital signal processing. This is useful in virtual meeting environments where each participant's voice has a constant localisation with respect to the video presentation irrespective of the listener's head rotation.

Ambisonics

Ambisonic systems are different from transaural and binaural schemes as they are not based on the human hearing system but on mathematical relationships describing how sounds originate from different locations combine acoustically. An ambisonic recording consists of three or four components: the first one describes the sound present in the left-right axis (the X-component), the second




Figure 4—By processing the audio presented to the user it is possible to create the impression of sounds appearing at locations around the user. Each voice is processed so that it appears to originate from a different location, giving clarity and a sense of immersion

describes the sound present in the front-back axis (the Y-component) and the third optional component describes the sound present in the up-down axis (the Z-component). An additional W-component describes the overall non-directional sound magnitude. The four signals W, X, Y and Z (known as a *B-format composite signal*) characterise the spatial audio information and are a very efficient mechanism for transmitting spatial audio over a network. At the recipient's location, the B-format signal is fed into an ambisonic decoder along with the positions of the loudspeakers and is decoded to give an accurate three-dimensional sound. As knowledge of the loudspeaker locations is only needed at the decoding stage, the same signal can be sent to several users, each with a different loudspeaker layout or different numbers of loudspeakers. Figure 4 shows a single user in both a transaural (blue loudspeakers) and an ambisonic (red loudspeakers) system. The ambisonic signal is decoded into four loudspeakers and so in this example there is no up-down information. Sounds can be made to appear as if they come from any location in the room—here there are six other users with each voice located at a different position around the user.

The main disadvantage of an ambisonic system is that for the best three-dimensional effect the user needs to be located precisely at the correct listening position. In environments where several users are located together, only one or two of the users will experience the correct spatial sound. At BT Laboratories, warping algorithms have been developed that give any user in a room the correct audio spatialisation, enabling all of the users to be fully involved in the communication experience.

Panning

The simplest method of producing a spatial audio signal is to position several loudspeakers around a room and pan sound between the loudspeakers to give an apparent location in space. The panning used can be either simple linear panning or a more complex non-linear panning technique. While simple, the main disadvantage of panning techniques compared to ambisonics is that the location of the loudspeakers needs to be known before the signal can be encoded. By comparison, an ambisonic system generates a spatial audio signal suitable for decoding and playback in any environment.

The BT/ARC VisionDome is an immersive environment comprising a large plastic dome with a large



Figure 5—A group of users inside the VisionDome. The users are experiencing an educational application, where they are travelling through time to experience the interior of an Egyptian pyramid

concave projection screen mounted on the internal surface. The dome can hold up to 15 users. A video projector is mounted in the centre of the dome with special lenses to cover the full field of vision. The VisionDome can be used for teleconferencing, telemedicine, entertainment, remote control (controlling a satellite in orbit, for example) etc (see Figure 5). The main difference between the VisionDome and the SmartSpace chair is that the VisionDome is intended to be a multi-user interface. This requires both audio and video systems that are able to provide immersion for more than one person at once.

Within the dome, there are nine loudspeaker channels: front-left, front-right, screen-left, screen-right, screen-centre, rear-left, rear-right, rear-centre and a sub-woofer. The screen channel speakers are placed behind the curved projection screen, with suitable modification to the audio to account for the effect the screen has on the sound waves travelling through it. Each channel has a separate feed from the audio spatialisation system (which can run three-dimensional panning and three-dimensional ambisonic-decoding programs).

The input audio signals come from a variety of sources such as film sound tracks (in ambisonic B-format), video

players, computer systems, audio samplers (multi-channel), telephones (mono) and microphones (mono). For a film sound track the audio is stored on a hard-disk recorder and synchronised to the video with time codes. With real-time computer-generated images the audio is triggered from a large library of sounds stored in the sampler and the computer controls where in space the sound is to be positioned. Sounds can be located in one position (such as conference delegate) or they can be moved around (such as an aeroplane or moving data object).

Because the VisionDome is a hard plastic structure and because it has a concave screen, certain acoustical challenges exist that needed to be overcome in order to be able to produce realistic spatial audio. When the dome is unfurnished and untreated, a 'whispering gallery' effect is obtained owing to the acoustic reflections. This problem is overcome by treating the area behind the screen with acoustic tiles and the other areas with acoustic absorbing materials.

The *network spatial audio server* (NSAS) provides a network-based shared meeting environment with high-quality spatialised audio. Each terminal consists of a networked PC connected to two centrally based computers, one acting as a spatial audio server and the other to control

the virtual world. Each user's screen shows a view of a computer-generated room with doors to other rooms (as shown in Figure 6). The user is represented by a computer-animation known as an *avatar*. The audio output from each user is picked up by a microphone and fed into the spatial audio server. The server calculates a spatial audio signal for each of the users, which includes all of the sounds the user can hear along with their direction. Using the system it is possible to have meetings with other people in virtual rooms and to move from room to room, meeting with different people in each room. The NSAS system includes room acoustic modelling to provide realistic models of the room's reverberation and the effect of sound travelling through open doorways.

The spatial audio server is capable of delivering a range of spatial audio formats depending on the individual's equipment set-up. For example, a desk-top system as shown in Figure 6 uses a transaural audio system and a large-screen system (as shown in Figure 7) with life-sized avatars uses an ambisonic system.

Conclusions

Communication systems are now poised to move from the current

simple 'voice from a handset' towards a fully immersive experience. There are exciting opportunities to combine spatial audio with wrap-around video screens to create a new generation of entertainment, computing and communications interfaces. One aim is to make the communication experience as natural as possible even to the point of making users unaware of the equipment around them.

This article has looked at three different example applications. The SmartSpace chair represents the office of the future, with a large wraparound video screen and a transaural audio system; the combination of wrap-around video and audio gives the user a sense of immersion and aids the communication process. The VisionDome is a large dome that can house a number of users with a video screen filling the users' field of vision. A range of audio techniques can be used in the VisionDome such as ambisonics and panning to generate a three-dimensional sound field. The SmartSpace chair and the VisionDome can be used for applications such as teleconferencing, telemedicine, data visualisation, education and entertainment. A slightly different example application is the network spatial audio server, which is used as way of meeting people in a virtual world. Spatial audio is used to aid navigation

Figure 6—NSAS desk-top terminal with a dummy seated in the user's seat. The system shown here uses a two-loudspeaker transaural system





Figure 7—NSAS full-size terminal with a dummy standing in the user's position. The system shown here uses an eight-loudspeaker ambisonic system



through the virtual world and to improve intelligibility when talking to a number of other users at the same time. The three examples clearly demonstrate the benefits of using of spatial audio in communication tasks.

Audio is a key part of any communication process and this article has considered some of the technology necessary to deliver realistic spatial sound.

Further Reading

- RIMELL, A.; and HOLLIER, M. P. Reproduction of Spatialised Audio in Immersive Environments with Non-ideal Acoustic Conditions. Presented at 103rd convention of the Audio Engineering Society 1997. Preprint No 4543.
- RIMELL, A.; and HOLLIER, M. P. Design and Implementation of 3-Dimensional Spatial Audio for Immersive Environments. *Proc. Institute Of Acoustics: Reproduced Sound* 13, 1997, 19(6), pp. 87-94.
- HOLLIER, M. P.; RIMELL, A.; and BURRASTON, D. Spatial Audio Technology for Telepresence. *BT Technol. J.*, Oct. 1997, 15(4), pp. 33-41.
- BURRASTON, D.; HOLLIER, M. P.; and HAWKSFORD, M. O. Limitations of Dynamically Controlling the Listening Position in a 3D Ambisonic Environment. *Journal of the Audio Engineering Society* (Abstracts), 1997, 45(5), p. 413, Preprint No. 4460.
- BEGAULT, D. 3-D Sound for Virtual Reality and Multimedia. Academic Press, 1994, ISBN 0-12-084735-3.
- HAWKSFORD, M. O. Letter on Sound Localisation Using HRTFs. *Electronics World and Wireless World*, Sept. 1996.
- MOLLER, H. Fundamentals of Binaural Technology. *Applied Acoustics*, 1992, 36, pp. 171-218.
- COOPER, D.; and BAUCK, J. Prospects for Transaural Recording. *Journal of the Audio Engineering Society*, 1989, 37(1/2), pp. 3-19.
- GERZON, M. A. General Metatheory of Auditory Localisation. Presented at the 92nd convention of the Audio Engineering Society, 1992, Preprint No. 3261.
- GERZON, M. A. Surround Sound Psychoacoustics. *Wireless World*, 1974, 80, pp. 483-486.
- MALHAM, D., and MYATT, A. 3-D Sound Spatialisation Using Ambisonic Techniques. *Computer Music Journal*, 1995, 19(4), pp. 58-70.
- GERZON, M. A. The Design of Precisely Coincident Microphone Arrays for Stereo and Surround Sound. Presented at the 50th convention of the Audio Engineering Society, 1975.
- WALKER, G.; TRAILL, D.; HINDS, M.; COE, A.; and POLAINE, M. VisionDome: A Collaborative Virtual Environment. *Br. Telecommun. Eng.*, Oct. 1996, 16, p. 217. Also at: http://www.labs.bt.com/people/walkergr/IBTE_VisionDome/index.htm
- KLEINER, M.; DALENBACK, B.; and SVENSSON, P. Auralization—An Overview. *Journal of the Audio Engineering Society*, 1993, 41(11), pp. 861-875.


Biographies



Andrew Rimell
BT Networks and System

Andrew Rimell completed an M.Sc. and Ph.D. at the University of Essex in 1993 and 1996, respectively. His Ph.D. proposed a method of reducing the perceived crossover-induced errors in a multi-driver loudspeaker system through using a psychoacoustically-motivated

error-concealment algorithm. In 1996, he was awarded an IEE Premium Award for his work in modelling room acoustics. He works at BT Laboratories in the Advanced Perception group. His work includes research into spatialised audio for telepresence and the use of psychoacoustic models in the quality assessment of telecommunications systems. He is an active member of the IEE Professional Group Committee E-10 (Circuits and Systems). He is also an active member of the AES and has presented numerous conference papers. His research interests include spatial audio, perceptual modelling, loudspeakers and signal processing.



Mike Hollier
BT Networks and Systems

Mike Hollier obtained a B.Eng. degree in Mechanical Engineering from Plymouth Polytechnic in 1987. After working as a hi-fi loudspeaker designer, he joined BT Laboratories to work on a variety of acoustic projects including the design of a novel noise-cancelling handset. He obtained an IOA diploma in acoustics in 1989, receiving an industry award for his research into the vibrational behaviour of light structures. From 1990, his work included the research of objective methods capable of predicting the performance of non-linear audio systems. He developed an analysis based on a model of human auditory perception, receiving his Ph.D. from the University of Essex in 1995 together with the BT Postgraduate publication prize. Since January 1996, he has led the Multi-Modal Perception group with projects on auditory, visual and multi-sensory perception. His team also works on spatial audio for conferencing and virtual environments. He is a visiting lecturer to the University of Essex, where he is a fellow of the Department of Electronic Systems Engineering, a Chartered Engineer and a member of the IOA and AES.

A Wearable Communication Space

As information and instantaneous access to it becomes increasingly important in the modern world, new technologies such as wearable computers are evolving to provide constant access to computing and communications resources. Using traditional virtual reality techniques, a wearable computer has been developed to provide access to such facilities. The wearable computer enhances communications through the use of spatialised three-dimensional graphics and audio, allowing natural head movements to be used to attend to the virtual representations much as they would be in the real world.

Introduction

As information continues to be an increasingly important commodity, the need is growing for more constant and immediate methods to access it. This, combined with the miniaturisation and integration of computing and communication devices, opens up the possibilities for the development of portable devices to fulfil this role. Such devices should allow instant access to both computer information and communications systems wherever the user is located. One of the major challenges in this area is the integration of natural communication into these devices, and it is this aspect that this article addresses.

Wearable computers are the most recent generation of portable machines, miniature versions of conventional personal computers which are worn on the body. In general, a wearable computer may be defined as a computer that is subsumed into the personal space of the user, controlled by the wearer and has both operational and interactional constancy; that is, always on and always accessible¹. Wearable computers are typically composed of a belt or backpack PC, see-through or see-around head-mounted display (HMD), wireless communications hardware and an input device such as touchpad or chording keyboard. This configuration has been demonstrated in several real-world applications including aircraft maintenance², navigational assistance³ and vehicle mechanics⁴. Such applications, where the user benefits from the virtual presence of an expert, have been found to improve task performance significantly^{5,6},

reducing task time by half in the case of vehicle inspection⁴.

Current wearable applications have, however, involved connections only between one local and one remote user. The problem currently being addressed is how the computing power of the wearable computer can be used to support collaboration between several people. In particular, the following issues are being explored:

- how wearable computers can be used to allow natural interactions with both data and people in a virtual world without impairing the ability to interact with the real world;
- what visual and audio enhancements can be used to aid such interactions; and
- how remote users can be represented in a wearable-computing environment.

These issues are becoming increasingly important as telephones incorporate more computing power and portable computers become more like telephones. However, a key question is whether it is necessary to use visual and audio enhancements in collaborative spaces—is computer-mediated communication needed when a conference telephone call may be just as effective?

The next section reviews relevant related research from the teleconferencing and collaborative virtual environment (CVE) fields. The article goes on to describe how communications and computing facilities can be combined in a wearable form to create a collaborative wearable communication space. Prototype

initial-user experiences and possible application areas are described.

Background

Previous research on the roles of audio and visual cues in teleconferencing has produced mixed results. Many experiments have been conducted comparing face-to-face, audio and video, and audio-only communications as summarised by Sellen⁷. While people generally do not prefer the audio-only scenario, they are often able to perform tasks as effectively or almost as effectively as when using face-to-face or video communication. While Williams⁸ finds that face-to-face interaction is no better than speech-only communication for cognitive problem-solving tasks, Chapanis⁹ finds that visual cues were important in tasks requiring negotiation. Even though the outcome may be the same, the process of communication is affected by the presence or absence of visual cues¹⁰, although not for managing turn taking¹¹. However, there is strong evidence that video transmits social cues and affective information, establishing 'social presence'¹², although not as effectively as face-to-face interaction¹³. When users attempt non-verbal communication in a video-conferencing environment their gestures must be wildly exaggerated to be recognised as the equivalent face-to-face gestures¹³. In general, the usefulness of video for transmitting non-verbal cues may be overestimated and video may be better used to show the communication availability of others or views of shared workspaces¹¹.

Based on these results, and the fact that speech is the critical medium in teleconferencing experiments¹¹, audio alone should be suitable for a shared communication space. An example of this, Thunderwire¹⁴, was a purely audio system which allowed high-quality audio conferencing between multiple participants at the flick of a switch. In a three-month trial, Hindus et al.

found that audio can be sufficient for a useable communication space and that Thunderwire afforded a social space for its users. However, several major problems were observed:

- users were not easily able to tell who else was within the space, and
- users were not able to use visual cues to determine other's willingness to interact.

In addition, Thunderwire was rarely used by more than two or three users at once. With more users it became increasingly difficult to discriminate between speakers and there was a higher incidence of speaker overlap and interruptions. These problems are typical of audio-only spaces and suggest that while audio may be useful for small group interactions, it becomes less useable when more people are present.

These shortcomings can be overcome through the use of visual and spatial cues. For example, the Passepartout¹⁵ enhanced desktop conferencing tool includes visual representation of a conference table alongside shared documents and text chat facility. The conference table includes icons of those people within the conference and simple cues, such as microphone on/off, which allow a participant's activity or interest to be inferred. In face-to-face interaction, speech, gesture, body language and other non-verbal cues combine to show attention and interest. However, the absence of spatial cues in most video-conferencing systems means that users often find it difficult to know when people are paying attention to them, to hold side conversations, and to establish eye contact¹⁶. This may explain the similarity in results between audio only and video-enhanced teleconferencing conditions, and the difference they both produce from face-to-face results.

Virtual reality can provide an alternative communication medium with groups of people being able to

share a virtual world, typically accessible via desktop computers connected to the Internet. BT has demonstrated virtual conferencing in which many users can be represented as life-like virtual 'avatars' of themselves within a virtual room¹⁷. In virtual environments, spatial cues can be combined with visual and audio cues in a natural way to aid communication¹⁸. The well-known 'cocktail-party' effect shows that people can easily monitor several spatialised audio streams at once, selectively focusing on those of interest¹⁹. In virtual environments, even a simple avatar representation combined with a spatial audio model of those in the environment enables users to discriminate between multiple speakers²⁰. Spatialised interactions are particularly valuable, enabling large numbers of people to inhabit the same virtual environment and interact in a way impossible in traditional video or audio conferencing²¹.

The remainder of this article describes how techniques from collaborative virtual environments can be used to develop a wearable communication interface.

A Wearable Communication Space

While considerable work has been conducted on the development of the use of immersive virtual environments for collaboration, there has been almost no research on collaboration in augmented reality settings. Billinghurst et al. found that users located in the same room prefer and perform better using augmented reality over a fully immersive interface²². Similar results have also been reported by Schmalstieg et al.²³. However, practical uses of wearable computers for communications dictate that the users are likely to be in different locations, and little research has been undertaken for such scenarios. An earlier paper describes how wearable computers




Figure 1(a) – Head-stabilised information display

are ideally suited for three-dimensional computer-supported collaborative working (CSCW) because they allow enhancement of real-world tasks and seamless interaction between the real world and virtual image overlays²⁴.

The results in the previous section suggest that an ideal wearable communication space should have three elements:

- high-quality audio communication,
- visual representations of the collaborators, and
- an underlying spatial model for mediating interactions.

One of the most important aspects of creating these communication interfaces is the visual and audio presentation of information. Most current wearable computers use see-through or see-around monoscopic head-mounted displays. With these displays, visual information can be presented in three ways:

- **Head-stabilised** Information is fixed to the user's viewpoint and does not change as the user moves in the real world.
- **Body-stabilised** Information is fixed relative to the user's body position and varies as the user's head is moved, but not as he/she changes position. This requires the user's head orientation to be tracked.
- **World-stabilised** The virtual information is superimposed on the real world, and the user can navigate around it much as he/




Figure 1(b) – One degree of freedom

she would move around real environments. This requires both the user's head position and orientation to be tracked.

Body- and world-stabilised information display is attractive for a number of reasons. As Reichlen²⁵ demonstrates, a body-stabilised information space can overcome the resolution limitations of head-mounted displays. In his work, a user wears a head-mounted display while seated on a swivel chair. By tracking head orientation, the user experiences a hemispherical information surround—in effect a 'hundred million pixel display'. Users can locate information more rapidly when using body-stabilised information displays than when using the more traditional head-stabilised versions²⁶. World-stabilised information presentation enables annotation of the real world with context-dependent data, creating information enriched environments²⁷. This increases the intuitiveness of real-world tasks. Despite these advantages, most wearables use only head-stabilised information display.

This article describes the use of the simplest form of body-stabilised display; one which uses one degree of orientation to give the user the impression he or she is surrounded by a virtual cylinder of visual and auditory information. Figure 1 contrasts this with the traditional head-stabilised wearable information presentation.

Body-stabilised display

In this case, head motion is tracked about the y (yaw) axis to change the user's view of the information space. Using only one degree of freedom has a number of advantages:


- users cannot become easily disoriented,
- no additional input devices are needed, and
- it is natural to use since most head and body motion is about the vertical axis.

When using a head-mounted display to navigate a cylindrical body-stabilised space, only the portion of the 'information space' in its field of view can be seen. The rest of the space can be viewed in two ways: by rotating the information space about the user's head, or by tracking the user's head as he or she looks about the space. The first requires no additional hardware and can be done by mapping mouse, switch or voice input to direction and angle of rotation, while the second requires a simple one-degree-of-freedom tracker. These minimal hardware requirements make cylindrical spatial information displays particularly attractive.

With this display configuration, remote collaborators can appear as virtual avatars distributed about the user (Figure 2). As they speak, their audio streams can be spatialised in real time so that they appear to emit from the corresponding avatar.

Just as in face-to-face collaboration, users can turn to face whoever they want to talk to while still being aware of any other virtual or real conversations which may be taking

Figure 2 – A spatial conferencing space.






Figure 3—The wearable interface

place. Since the displays are see-through or see-around, the user can also see the real world at the same time, enabling others in the virtual environment to help with real-world tasks. These remote users may also be using wearable computers and head-mounted displays or could be interacting through a desktop workstation.

Implementation

This research is initially focused on collaboration between a single wearable-computer user and several desktop PC users. The aim is to develop software to support medium-sized meetings (five–six people) in a manner that is both natural and intuitive to use.

The wearable computer used is a custom-built 586 PC with 20 Mbyte of RAM running Windows '95. A hand-held Logitech wireless radio trackball with three buttons is used as the primary input device. The display is a pair of Virtual i-O iglasses converted into a monoscopic display by the removal of the left eyepiece. The Virtual i-O head-mounted display can either be used in see-through or occluded mode, has a resolution of 262 by 230 pixels and a 30-degree field of view. The iglasses have stereo headphones, an inertial and gyroscopic three-degrees-of-freedom orientation tracker. A

BreezeCom wireless local area network (LAN) is used to give 2 Mbit/s access to computer networks up to 150 metres from a base station. The wearable computer also has a 'soundblaster' compatible sound board with head-mounted microphone. Figure 3 shows a user wearing the display and computer. The desktop PCs are standard machines with network connectivity and sound capability.

Wearable interface

The wearable computer has no graphics acceleration hardware so the interface was deliberately kept simple. The conferencing space runs as a full-screen application that is initially blank until remote users connect. When users join the conferencing space they are represented by a simple avatar consisting of a 128 × 128 pixel texture mapped picture of themselves. Although the resolution of the images is crude it is sufficient to identify who the speakers are. The wearable user's head is tracked so he or she can simply turn to face particular speakers. As the user faces different speakers the volume of the speaker's speech changes through the use of three-dimensional sound spatialisation. Users can navigate through the virtual space: by rolling a trackball forwards or backwards their viewpoint is moved forwards or backwards along the direction they are looking. Since the virtual images are superimposed on the real world, when users roll the trackball it appears to them as though they are

moving the virtual space around them, rather than navigating through the space. Figure 4 shows the wearable interface from the wearable user's perspective.

Desktop interface

Users at a desktop workstation interact with the virtual conferencing space through a similar interface as the wearable user. In this case, however, it is a windows application on the desktop and the space is navigated by using a mouse. The virtual space is rotated using mouse movements when the left mouse button is held down; mouse movements with the left button released allow the user to move backwards and forwards in the space. Mapping avatar orientation to mouse movements means that the desktop interface is not quite as intuitive as the wearable interface. Users at the desktop machine wear head-mounted microphones to talk in the conferencing space.

Initial User Experiences

The effect of using spatialised audio and visual representations on communication is being evaluated by conducting user trials. Pre-trial results show that users are able to discriminate easily between three speakers when their audio streams are spatialised, but not when non-spatialised audio is used. Users also prefer seeing visual representations of the others in the environment as opposed to just hearing their speech. They found that they could continue doing real-world tasks while talking to others in the conferencing space and that they could move the virtual conferencing space around with the trackball so that its elements would not block critical portions of the user's view of the real world.

However, as more users connect to the conferencing space the need to spatialise multiple audio streams puts a severe load on the processor;

Figure 4—The user's perspective



slowing down the graphics and head tracking. This makes it difficult for the wearable user to conference with more than two or three people simultaneously. This problem will be reduced as faster processors and hardware support for three-dimensional graphics become available for wearable computers. Spatial culling of the audio streams could also be used to overcome this limitation, through the coagulation of selected streams to a single spatial location, or by removing some of the audio altogether.

Discussion


This article has described a prototype wearable communication space that uses spatial visual and audio cues to enhance communication between remote groups of people. The interface shows what is possible when computing and communications facilities are coupled together on a wearable platform. Preliminary results have found that users prefer using both the audio and visual cues together and that spatialised audio makes it easier for users to discriminate between speakers. Although only minimal avatar representations are used, users found that they enhanced the ability to communicate effectively. Formal studies are being conducted to confirm these results and evaluate the effect of spatial cues on communication patterns. Future work is planned to investigate how the presence of spatialised real-time video can further enhance communication. Live video texture mapping will be incorporated into the interface, enabling users to see each other as they speak. This will also allow users to send views of their workspace to each other, improving collaboration on real-world tasks.

A wearable communication space can be used to support numerous collaborative applications in which some participants are either not sitting at desks or need mobility. A

shared virtual environment facilitates audio-visual communications for groups of people with the added potential for embedded graphical, textural or audio information. A particularly important trait is that in an 'augmented' communication space it is possible for users to form effective 'cognitive maps', by associating the annotated information with physical objects within their surroundings. To date we have demonstrated a 'body stabilised' system, in which the communication space is located relative to the user themselves. With the ability to position explicitly the communication space, and objects within it, relative to the user's real-world location, an exciting range of possibilities becomes possible. A user could, for example, choose to view avatars of conference participants or shared information overlaid and 'attached' to a physical noticeboard. This represents a powerful vision of conferencing for all platforms; with remote video-conference participants not in separate windows on a screen but spread around the user's environment, positioned in space where the user prefers.

Perhaps the most significant limitation of wearable computers is the current intrusive displays which impose unacceptable social constraints within many applications. For example, it is difficult to be taken seriously during face-to-face conversations when wearing the

Figure 5—The 'Eyeglass' display



display shown in Figure 3. However a commercial alternative that combines a colour VGA resolution LCD screen with in a pair of spectacles is close to reality. The 'Eyeglass Display' from MicroOptical, as shown in Figure 5, will undoubtedly accelerate the use and ubiquity of wearable computers (<http://www.microopticalcorp.com/>).

Acknowledgements

We would like to thank our colleagues within the Applied Research and Technology Department of BT and Nick Dyer for producing the renderings used in some of the figures.

References

- 1 MANN, S. Smart Clothing: The 'Wearable Computer' and WearCam. *Personal Technologies*, 1(1) March 1997, Springer-Verlag.
- 2 ESPISITO, C. Wearable Computers: Field-Test Results and System Design Guidelines. Proceedings Interact '97, 14–18 July 1997, Sydney, Australia.
- 3 FEINER, S.; MACINTYRE, B; and HILLERER, T. A Touring Machine: Prototyping 3D Mobile Augmented Reality Systems for Exploring the Urban Environment. Proceedings of the International Symposium on Wearable Computers, Cambridge, MA, 13–14, October 1997, pp. 74–81, IEEE Press, Los Alamitos.
- 4 BASS, L.; KASABACH, C.; MARTIN, R.; SIEWIOREK, D.; SMAILAGIC, A.; and STIVORIC, J. The Design of a Wearable Computer. Proceedings of CHI '97, Atlanta, Georgia. March 1997, pp. 139–146, ACM Press, New York.
- 5 SIEGAL, J.; KRAUT, R.; JOHN, B.; and CARLEY, K. An Empirical

- Study of Collaborative Wearable Computer Systems. Proceedings of CHI 95 Conference Companion, 7–11 May 1995, Denver, Colorado, pp. 312–313, ACM Press, New York,
- 6 KRAUT, R.; MILLER, M.; and SIEGAL, J. Collaboration in Performance of Physical Tasks: Effects on Outcomes and Communication. Proceedings of CSCW '96, 16–20 Nov. 1996, Cambridge MA., ACM Press, New York.
- 7 SELLEN, A. Remote Conversations: The Effects of Mediating Talk with Technology. *Human Computer Interaction*, 1995, **10**(4), pp. 401–444.
- 8 WILLIAMS, E. Experimental Comparisons of Face-to-Face and Mediated Communication. *Psychological Bulletin*, **16**, 1997, pp. 963–976.
- 9 CHAPANIS, A. Interactive Human Communication. *Scientific American*, **232**, 1975, pp. 36–42.
- 10 O'MALLEY, C.; LANGTON, S.; ANDERSON, A.; DOHERTY-SNEDDON, G.; and BRUCE, V. Comparison of Face-to-Face and Video-Mediated Interaction. *Interacting with Computers*, **8**(2), 1996, pp. 177–192.
- 11 WHITTAKER, S. Rethinking Video as a Technology for Interpersonal Communications: *Theory and Design Implications*. Academic Press Limited, 1995.
- 12 WHITTAKER, S.; AND O'CONNAILL, B. The Role of Vision in Face-to-Face and Mediated Communication. *Video-Mediated Communication*, 1997, pp. 23–49. Erlbaum Associates, New Jersey.
- 13 HEATH, C.; and LUFF, P. Disembodied Conduct: Communication Through Video in a Multimedia Environment. Proceedings of CHI '91 Human Factors in Computing Systems, 1991, pp. 99–103, ACM Press, New York.
- 14 HINDUS, D.; ACKERMAN, M.; MAINWARING, S.; and STARR, B. Thunderwire: A Field Study of an Audio-Only Media Space. Proceedings of CSCW '96, 16–20 Nov. 1996, Cambridge MA., ACM Press, New York.
- 15 RUSS, M., Desktop Conversations—the Future of Multimedia Conferencing, *BT Technol. J.*, Oct. 1997, **14**(4), pp. 42–50.
- 16 SELLEN, A. Speech Patterns in Video-Mediated Conversations. Proceedings CHI '92, 3–7 May 1992, pp. 49–59, ACM Press, New York.
- 17 MORTLOCK, A.; MACHIN, D.; MCCONNELL, S.; and SHEPPARD, P. Virtual Conferencing, *BT Technol. J.*, Oct. 1997, **14**(4), pp. 120–129.
- 18 BENFORD, S.; and FAHLEN, L. A Spatial Model of Interaction in Virtual Environments. Proceedings of Third European Conference on Computer Supported Cooperative Work (ECSCW'93), Milan, Italy, Sept. 1993.
- 19 BREGMAN, A. Auditory Scene Analysis: The Perceptual Organization of Sound, 1990, MIT Press
- 20 NAKANISHI, H.; YOSHIDA, C.; NISHIMURA, T.; and ISHIDA, T. FreeWalk: Supporting Casual Meetings in a Network. Proceedings of CSCW '96, 16–20 Nov. 1996, Cambridge MA., pp. 308–314, ACM Press, New York.
- 21 BENFORD, S.; GREENHALGH, C.; and LLOYD, D. Crowded Collaborative Virtual Environments. Proceedings of CHI '97, Atlanta, Georgia, March 1997, pp. 59–66, ACM Press, New York.
- 22 BILLINGHURST, M.; WEGHORST, S.; and FURNESS, T. Shared Space: Collaborative Augmented Reality. Proceedings CVE '96 Workshop, 19–20 Sept. 1996, Nottingham, UK.
- 23 SCHMALSTEIG, D.; FUHRMANN, A.; SZALAVARI, Z.; and GERVAUTZ, M. Studierstube—An Environment for Collaboration in Augmented Reality. Proceedings CVE '96 Workshop, 19–20 Sept. 1996, Nottingham, UK.
- 24 BILLINGHURST, M.; WEGHORST, S.; and FURNESS, T. Wearable Computers for Three-Dimensional CSCW. Proceedings of the International Symposium on Wearable Computers, Cambridge, MA., 13–14 Oct. 1997. pp. 39–46. IEEE Press.
- 25 REICHLEN, B. SparcChair: One Hundred Million Pixel Display. Proceedings IEEE VRAIS '93. Seattle WA, 18–22 Sept. 1993, pp. 300–307, IEEE Press.
- 26 BILLINGHURST, M.; BOWSKILL, J.; DYER, N.; and MORPHETT, J. An Evaluation of Wearable Information Spaces. Proceedings of the IEEE Virtual Reality Annual International Symposium (VRAIS 98), Atlanta, 14–18 March 1998.
- 27 REKIMOTO, J.; and NAGAO, K. The World through the Computer: Computer Augmented Interaction with Real World Environments. Proceedings of User Interface Software and Technology '95 (UIST '95), November 1995, pp. 29–36, ACM Press, New York.

Biographies



Mark Billinghamurst
University of
Washington

Mark Billinghamurst is a researcher at the Human Interface Technology Laboratory, University of Washington, where he co-manages wearable computing work. He has published over 30 papers in the areas of virtual reality (VR) and advanced interface design. Past projects include voice and gestural interfaces, evaluation of VR interaction techniques, intelligent virtual interfaces, and collaborative augmented reality environments. His current work involves using VR techniques to develop interfaces for wearable computers, particularly interfaces that support collaboration on a wearable platform. He has a M.Phil. in Applied Mathematics from the University of Waikato, New Zealand, and is in the final year of completing a Ph.D. in Electrical Engineering from the University of Washington, Seattle. More information can be found at:
<http://www.hitl.washington.edu/people/grof>



Jerry Bowskill
BT Networks and
Systems

Jerry Bowskill joined the Advanced Perception Unit within the Applied Research and Technology Department at BT Laboratories in 1996. Before joining BT, he worked as a Research Officer within the Department of Electrical Engineering at the University of Brighton. In 1990, he was awarded an honours degree in Microelectronics and Information Processing. After graduating, he started his research career at the University of Brighton as a research assistant working on machine-vision systems. During the past three years he has been researching 'enhanced reality'. He was awarded a Ph.D. in 1994 and has published widely in the areas of concurrent engineering, machine vision, virtual and enhanced reality.



Jason Morphet
BT Networks and
Systems

Jason Morphet completed his B.Sc. in Software Engineering at Suffolk College (associate college of the University of East Anglia) following an HND in the same subject. He is currently a registered Ph.D. student at the University of Nottingham. He joined BT as a summer student in June 1995, staying on to become a Research Fellow and subsequently a full-time employee. His interests include both the network and human interface issues of distributed virtual environments, and he was closely involved in Advanced Applications and Technology (AA&T) Portal development.

Interacting Demands for Multimedia Services

An explosion in the provision and use of multimedia services has been widely heralded. However, it is impossible to predict accurately the impact of these services. There are key issues which not only are critical to network investment commitments but are also largely unknown, subject to enormous errors, with potentially catastrophic financial implications. It is critical to understand all these interacting parameters. A market forecast modelling technique has been developed to assess the impact and scale of such issues.

This article is based on the paper presented by the author at the IBTE/FITCE colloquium on 'Multimedia and the Internet', London, 31 October 1997

Introduction

In this article, a modelling methodology will be discussed, which will show that to dimension telecommunications networks efficiently and cost-effectively it is important to take into account, over a number of years, the changes in the traffic generated and customer numbers for each individual service. If this is not done, there is a risk that the network will be under-dimensioned, necessitating expensive upgrades.

Furthermore, in order to avoid costly over-dimensioning of the network, it is necessary to take into account that the peaks of traffic for each service could occur at different times of the day. For instance, we might expect traffic from schools services to be greatest from 09.00 hours to 15.30 hours and games traffic to peak before 09.00 hours and after 15.30 hours when the children are at home. In this case, it would be wrong to dimension the network to cope with the sum of the maximum traffic level for schools, occurring say at 13.00 hours, and the maximum traffic level from games, occurring say at 19.00 hours. Rather we should find the maximum traffic level for the combined services. This might be at, say 16.00 hours, while some children are still at school and some have just arrived home and are keen to use their network games. In this way it should be possible to avoid unduly high initial capital outlay for networks which are over-dimensioned.

However, care must be taken as a network design which might appear to be correctly dimensioned in year 1 could be significantly under-dimensioned by, say, year 5, if customer numbers, service offerings, service usage and the subsequent traffic levels

grow rapidly. That is, it will be necessary to examine a number of 'what-if' scenarios to ascertain which investment and network dimensioning options are to be chosen, not only for today, but also into the future. In order to do this, it is necessary to understand the problem space and to create a model to assist the scenario exploration.

The Problem Space

Let us consider now the issues concerning telecommunications companies in their own domestic markets. The market can be split into several segments such as residential, business, education, health and so on. Each one of these market segments will have its own multitude of service offerings, with different numbers of customers, tariffs, usage levels and time-of-day traffic profiles. Further, all of these parameters will evolve with time. This creates a complex problem space of numerous interacting parameters. Consider Figure 1, which shows just four market segments, some example services and how they contribute to traffic levels.

Figure 1 – The creation of traffic from different market segments




Figure 2—Flow diagram of the model


Figure 1 shows that, just by considering traffic, we need to understand the traffic generation from each individual service, its time-of-day traffic profile and how the traffic generated will grow over the years. So, there is already a complex set of parameters to analyse.

However, this problem space could be further complicated. Each of these segments could be subdivided. For example, we may wish to segment the residential market according to demographic groupings relating to wealth, such as A, B, C, D and E, and segment again into family or non-family. As another example, the business segment could be grouped according to the Standard Industrial Classification and then split again into small, medium and large enterprises. The education segment could be separated into primary, secondary and tertiary education which could be sub-divided again into state, grant maintained, independent, technology colleges, sixth-form colleges and so on. The community segment could be divided according to private residential care, hospital care, and split again into passive monitoring, interactive monitoring and so on.

This already shows that the problem space is in reality considerably more complex than illustrated in Figure 1. However, this is still just a subset of the problem space. There are more market segments than the four illustrated here. We could include, for instance, mobile communications, interconnect with other operators, international communications and so on.

But do not forget that we have been considering only traffic. This is not the only quantity that will vary with time (and time too, can be defined as hours of the day, days of the week, months and years). We need to remember that customer numbers, tariffing structures, revenues, etc. will also vary.

With such an enormously complex problem space, one might be tempted to ask, 'why should we bother to model this?', 'won't the errors be too



large?', 'how can we have confidence that our model gives realistic scenarios?' We shall consider these issues in the next section. However, it should be remembered that if the forecast of the explosion in multimedia services occurs, then telecommunications companies which fail to address how this problem space impacts their business face potentially catastrophic financial loss. Consequently, it is extremely important that the model should imitate reality closely. This means that the model needs to include as many of the real-life variables as possible.

Creating a Framework for the Model

We have already identified that the problem space is enormously complex. To model such a problem space poses tremendous problems in validating the outputs. To do this we need to break the problem space into small, manageable jigsaw pieces for each market segment. These jigsaw pieces could be considered to be, for instance, customer migration or tariffing structures or service usage or time-of-day traffic profiles and so on. These individual pieces can be tested and then fitted together to form the overall model for that individual market segment.

By repeating this process for each individual market segment, we can then fit the segments together to form the overall picture for the telco.

Let us now consider modelling the market segments and proceed to assess how the jigsaw pieces should be formed. For each market segment,

the problem space can be represented by a flow diagram of interacting segments, as shown in Figure 2.

For each market segment, we can consider that a telco has an investment policy which will enable it to provide new services and networks to *Potential Customers for New Services*. However, some of these customers will be reluctant to change and will keep their existing services and their existing rental and call charges will contribute to the telco's revenue, as they have always done. Some of the *Potential Customers for New Services* will migrate to the new services and will become *Customers of New Services*. This will generate connection and rental charges from the new services which will contribute to the revenue. As the customers of new services use their new services, so traffic levels will begin to rise and pay-as-you-use charges will also contribute to the revenue generated. With time, increasing numbers of customers of new services and rising customer use of services will lead to increasing traffic. This may necessitate network upgrade. The initial investment policy and network upgrade both contribute to the costs incurred by the telco. Then revenue minus costs will give the profit to the telco.

Notice again that the modules in the flow diagram could be subdivided. For instance, in the traffic module, time-of-day service profiles and yearly variations could each be modelled. *Customers of New Services* could be categorised according to whether their service offering is basic or has new enhanced features (for example, higher

Figure 3—Stacking customer segments to calculate customer migration, service usage and traffic

bit-rate). Costs could include whole-life costs as well as initial capital costs. These costs could be sub-divided and apportioned to particular services.


The flow diagram shown in Figure 2 has been modelled using 'Powersim' which applies systems dynamics to assess time-varying, dynamic parameters to generate a 'what-if' model to explore scenarios. The resultant model is known as the *Integrated Networks and Services Model* (INMS).

Due to the complexity of the modelling procedure, for the remainder of this article, we shall consider just the migration of customers, service offerings (and subsequent enhancements), service usage, time-of-day traffic profiles and traffic evolution over the years. This will be performed for five key services: Internet, Video-on-Demand (VoD), Broadcast TV and Games for the residential segment, and primary school services (which in this case will be defined as an Internet-type service, for the education segment).

Structuring the Sub-Modules for the INMS

The flow diagrams governing the analysis (or 'jigsaw pieces'), of customer migration, service take-up and usage, time-of-day traffic profiles and traffic evolution will now be discussed.

As an example, the five key services mentioned previously will now be considered. It will be necessary to analyse the migration of customers in the residential and primary education segments to these new services in order to estimate the service usage and subsequent traffic levels. However, not all residential customers are the same. Further, there can be considerable differences in the requirements of primary schools. So, the residential segment will be sub-divided into five demographic groups according to wealth as A, B, C, D and E. Here, A represents the wealthiest group and E represents the poorest group. We shall subdivide these demographic groups into family and non-family, giving a



total of 10 groups in the residential segment. The primary education segment will be subdivided into three groups, state, grant-maintained and independent. Figure 3 shows how *Customers of Existing Service* can migrate to become *Customers of New Basic Service*, (where the new basic service is one of the five services we are studying, but in its most primitive form; for example, low bit-rate). These *Customers of Existing Service* and *Customers of New Basic Service* can, after the subsequent launch of a new enhanced version of the service (for example, higher bit-rate), migrate to become *Customers of New Enhanced Service*. Now, *Usage of New Basic Service* and *Usage of New Enhanced Service* will lead to the creation of total up- and downstream traffic. This traffic can be allocated a daily profile so that the hourly time-of-day traffic can be calculated. Analysis of the time-of-day traffic, (most especially, the peak hour), and the total up- and downstream


traffic shows whether the network is over- or under-dimensioned.

Note that the flow diagram shown in Figure 3, is applied to each of the individual customer groups shown, (10 for the resident segment and three for the primary-education segment). That is, in this example, the analysis in the flow diagram is performed 13 times. Note too that in the diagram, the residential segment 'Families' are represented by a white layer whereas 'Non-Families' are represented by a grey layer.

However, this stacking of customer migration, service usage and traffic for the 13 groupings has only been shown for one service (which is applied to the appropriate segment, either residential or education). We now need to stack all the services, Internet, Broadcast TV, VoD and Games, and Education, against the right market segment(s).

Figure 4 shows the stacking of the five services which each have a

Figure 4—Stacking service definition and usage



service definition for both its basic (initial) version and its enhanced version. These definitions are in terms of factors which affect traffic generation, such as bit-rate and the date when the service is first available. Customers will have an *Initial Service Usage* defined by the novelty factor and their unfamiliarity with this service. However, with time service usage trends will change and these are taken into account.

From Figures 3 and 4 we see that the model is stacking 10 of the (residential) demographic groupings, against four services and three of the (Education) demographic groupings against one service. That is, the model has 43 stacks. The INMS calculates the customer migration, service usage and traffic generation for all of these 43 individual stacks and integrates the results together. In the next section we shall consider some of the integrated results.

Integrated Results

Let us consider now one artificial example, of the traffic output and its five-year trends, in order to demonstrate some of the functionality of the model. Note, that this gives just one scenario, but, it is extremely important to model several scenarios to gain a good understanding of the implications to the network. In this example, the residential sector has been split into its constituent (demographic) parts and recombined to give hourly traffic levels for the four residential services of Broadcast TV, VoD, Internet and Games. Likewise, the primary-schools sector has been split and recombined to give the hourly traffic levels for the education service. Figures 5 and 6 show hourly traffic levels for year 1. Similarly, Figures 7 and 8 show hourly traffic levels for year 5. In each case the contribution from each individual service is shown, as is the total traffic from all the services. Note that Broadcast TV is absent from the year 1 curves, as it is assumed that this service has not yet been launched.




Figure 5 – Year 1 Traffic for All Services




Figure 6 – Year 1 Traffic for Internet and Education

Figure 5 shows that in year 1, VoD makes the most significant contribution to the total traffic. In this particular example, peak usage time for VoD occurs at approximately 23.00 hours. This is close to the total peak usage time for all the services which occurs at approximately 22.00 hours. Earlier in the day, between 05.00 hours and 09.00 hours, games usage makes the most significant contribution to the total traffic generated (although at a much lower level than the VoD traffic generated later in the day). Notice that VoD

and games make a much greater contribution to the total traffic than Internet and Primary Education. This is principally due to Internet and Primary Education operating at significantly lower bit-rates. Also, in this example, only a few primary schools have been connected to the service, as it has only just been launched.

Figure 6 shows an enlarged scale, with total traffic, VoD and Games removed, in order to illustrate the traffic profiles for Internet and Education. This shows that the




Figure 7—Year 5 traffic for all services




Figure 8—Year 5 traffic for Internet and education

contribution to traffic from education is indeed very small, (as modelled within this particular scenario).

Consider Figure 7 which shows the time-of-day traffic for each of the individual services and the total traffic in year 5. Notice that in just five years the total traffic level has grown by a factor of 30. Much of this growth is attributable to Broadcast TV now being carried across the network. Also, customer numbers in each of the services has been growing and all of the services have been

enhanced to operate at considerably higher bit-rates than in year 1. In this example, Broadcast TV is now the greatest generator of traffic, which peaks at about 21.00 hours. VoD is the second highest creator of traffic, which peaks at about 23.00 hours, as before. Again the total traffic peaks at about 22.00 hours

Figure 8 shows for year 5 the hourly traffic profiles of just games, Internet and education, on an enlarged scale. In this case, Internet traffic has grown to be much larger

than in year 1 and is now generating far more traffic than games, (which is the reverse of year 1). This means that we need to be aware that a service with relatively little traffic generation in its early years may have a growth in traffic which outstrips that of other services. That is, the network designer needs to have a strong knowledge of the implications of differential traffic growth from a variety of services.

Early in this article, it had been suggested that perhaps Education could have a significant contribution to the total daily traffic. In the particular scenario analysed, Education even in its busiest hour, makes a very small contribution to the total traffic. However, it must be remembered that secondary and tertiary education have not been included in this analysis. Further, if usage or service bit-rate increase significantly, then Education could dramatically affect the total-traffic profile. These are further scenarios which could be investigated.

The main conclusion from this particular scenario is that within five years, traffic generated is likely to grow by a factor of 30. This means that it is essential that in year 1 the network should not be dimensioned to cope with just year 1 traffic. Rather, in year 1 the network should be greatly over-dimensioned to cope with a rapidly increasing level of traffic. If this is not done then very large upgrade costs are envisaged.

Conclusions

A model has been developed to assess the effects of customer migration to new services, service usage, traffic generation and evolution, and the impact on the network. It has been created by breaking down the enormously complex and multi-dimensional problem space into smaller 'jigsaw' pieces which are developed and tested prior to fitting these pieces together to form a market segment sub-model. Different market segments are stacked together to create the entire model.

Many scenarios can be explored, of which only one artificial example for traffic generation has been illustrated. Note that, for a single simulation run of the model, many other parameters and their trends can also be output. Examples of the outputs include customer migration, revenue generation, costs, profit for each demographic sector and each service, as well as top-level overview scenarios.

The INMS is a valuable tool to aid the understanding of the enormous opportunities and risks faced by telecommunications companies in today's and tomorrow's multimedia world. This model offers the opportunity to avoid potentially catastrophic decisions relating to networks and services provision which could have disastrous financial implications.

Acknowledgements

The author wishes to thank Richard Nicol, Alan Steventon, Michael Lyons and Jo Bedward for reviewing this article.

Biography



Ann Matthews
BT Networks and
Systems

Ann Matthews joined BT in 1985 with a B.Sc. (Hons.) in Physics and Mathematics and an M.Sc. in Modern and Applied Optics, from the University of Reading. She designed and analysed theoretical and practical optical telecommunications systems until 1992. During this time, she obtained her Ph.D. in Novel Optical Design Theory from the University of Reading (in 1988). In 1992, she began business case modelling for proposed telecommunications networks and their services. She led a team in the development of the MANTIS model, which was a precursor to the INMS. MANTIS has been used in the analysis of several BT network proposals, and in several collaborative European projects. She initiated the development of the INSM in 1995, and now leads a team continuing the research. She also provides comparative competitor analysis on a consultancy basis within BT.

Creating the Corporate Soloist

Who plays the Stradivarius?

'Organisations learn only through individuals who learn. Individual learning does not guarantee organisational learning, but without it, no organisational learning occurs'

Peter Senge, author of *The Fifth Discipline – The Art and Practice of the Learning Organisation*, 1990

Introduction

The age of the individual has arrived with a vengeance. This idea is absolutely key to the notion of facilitating learning and creating knowledge across the organisation. No longer can organisations stay competitive if they simply rely on strategy, structure and processes to provide the market edge. These factors are important in helping to manage the control systems, but innovation, learning and knowledge will be managed and mastered by individuals, not faceless corporate bodies. A clear analogy for this can be seen with the idea of a musical soloist, the individual who is able to switch seamlessly between the role of team member and featured player in a band or orchestra.

If one considers the role of solo guitarists, violinists or drummers, they have to operate on a number of levels at any one time. They must:

- have the appropriate knowledge, skills and competencies to stand out from the band or orchestra, and be prepared to put their talents on public view;
- be able to adapt to constantly changing patterns and themes in the music;
- have the curiosity and desire to challenge themselves, and possibly the composer who wrote the original score;
- be prepared to reflect on, and learn from, each solo performance;
- have a strong relationship with the band, one built on trust, respect and freedom; and
- have a sense of consideration and allow others to display their proficiency.

The question that faces any band is how to keep an equitable balance. How can the skills and talents of the soloist be used to enhance the capability of the band, while not overshadowing everyone else? Correspondingly, how, in organisational terms, can the dynamic potential that lies within each of the people be used to enhance the capability of the business? In doing this, can the organisation offer an equitable balance, both helping individuals to realise their potential and stimulating new opportunities for the company?


Solo traits

This equitable balance is achieved by helping people to develop a clear sense of value—both the value that they contribute to the business, and the value that the business can offer with their personal development.

To help achieve this, six key traits need to be encouraged and developed in the people (Figure 1):

- *Compelling Future*—developing clear personal goals, and under-

Figure 1—Soloist six-trait model



This article reflects only the personal views of the author

Only by helping all individuals to develop a passion for their future will an animated sense of anticipation for the organisation be realised

standing what personal attributes will help to achieve them.

- *Change Mastery*—moving from a position where change is feared, to one where it is regarded as an invigorating and opportunistic process, even seeing change as a form of personal long-term security.
- *Continual Learning*—to understand that as an employee, one has become a marketable commodity, and as such, it is wise continually to re-invent and invest in one's asset base.
- *Challenging Ethos*—recognising that change and learning will always be about modifying paradigms and perceptions. To do this, one will always have to challenge the existing status quo and feel comfortable breaking with past conventions and habits.
- *Contract Alignment*—to be at ease with the change in the psychological contract with the organisation. Specifically, one must look at the relationship from a win/win perspective, where both parties have a sense of shared mutual benefit.
- *Considerate Nature*—recognising that it is fine to have personal goals, but their achievement will come through working with, and having interdependence with, others.

By no means are these six traits exclusive to the goal of creating individuals who are able to play a solo role. However, by focusing on these areas, the process of creating people who have a sense of mastery of their lives is somewhat easier, and more structured in its approach.

Trait 1—Compelling Future

If people are not excited about their future, then they are unlikely to be thrilled about the organisation's future. Without this personal stimulation, then what chance is there that

any degree of innovation, change or learning will take place? Only by helping all individuals to develop a passion for their future will an animated sense of anticipation for the organisation be realised. In creating this compelling future, individuals need to go through a number of introspective processes:

- *Knowledge of one's own values*—understand what personal drivers the individuals have at their disposal.
- *Find something special about oneself*—being clear on one special thing where they excel.
- *Time-line mentality*—being able to see themselves in a clear perspective against the past and future.

In building on each of these stages, individuals will have a clear, communicable and compelling sense of the future, for themselves and the organisation.

Knowledge of one's own values

One of the biggest downfalls that an individual can make is to create a personal compelling future that is built around another person or organisation's values and beliefs. All too often, people are fooled into thinking that their personal goal is to become a corporate leader; a qualified professional, or a respected practitioner; without really questioning where this belief has come from. Is it something that they truly desire, or is it just a manifestation of someone else's dreams they have been encouraged to accept?

In the past, the notion of a person working for one organisation (or trade) for life was often the norm. For example, individuals might join a company as trainees, then work their way slowly up the organisational ladder; desperately reaching for each new step, as though this was the only rationale for going to work. This often resulted in a situation where people never took the time to reflect on what

was important for them in life. As a model for managing a working life, the idea of focusing one's career on a single company or profession is dangerous. People need to think in terms of life-long learning, rather than a life-long employer. An individual must move from an organisational relationship founded in dependency to one based upon interdependency.

To create a personal and compelling future, individuals must first understand what it is they personally value in life. A recent study of exceptional companies led to the conclusion that it is more important to know who you are than where you are going, for where you are going will change as the world around you changes¹. Only once these personal values are understood can the next stage begin; that is, to find out what it is that makes the person unique.

Find something special about oneself

It is important to recognise that people are unique. As such the goal must be to bring out the range of diversity and dissimilarity, and so create variety and innovation across the organisation.

The successful organisation will seek to raise the level of diversity among its population. In working in this way, this individual uniqueness can be used in a productive way, and not be seen as something that the organisation has to subordinate or subsume. Imagine a culture that encourages normalised groups, stereotypes and people who are all prepared to act in the same way and follow the party line explicitly. This is one that might thrive in a static and stable environment, but will soon be overcome as the forces of change sweep through the market place. Alternately, imagine an organisation populated by complementary cultures, differing perspectives, and a greater appreciation of the varieties in life. This is one that will understand and adapt to the market as it transforms itself. Diversity, constructive conflict, turbulence and coinci-

dence are the seeds for learning and innovation. As such, it is imperative that people are encouraged to search this out, to learn how best to use it for themselves and the organisation.

The individual should consider to the extent to which they are truly open to differences, interested in others, and curious about new things. In support of this, the organisation needs to look inwards and consider its systems, processes and practices, to understand how they drive the behaviour of the people in the organisation. Do systems restrict diversity and tolerance? Do the processes encourage conformity, and is it accepted practice to prevent open disagreement? If so, what actions can be taken to change the embedded and underlying structures, and so move towards a more open and honest style of management?

Time-line mentality—bringing the future forward

In defining and creating a compelling future, one of the most dominant factors is the sense of time associated with the end position. The way that change is framed is always influenced by the perception of time, both the past and the future. The goal is for individuals to be able to map clearly their personal time lines and indicate where their goals fit on it.

The time line is a powerful tool that helps people to understand how they view the world. Some people look at their time line from a historic perspective, always reflecting on things that have (or might have) happened. Others tend to sit in the present, just focusing on today's issues without thinking about where they might be heading. To develop the sense of a compelling future, the idea is to operate out of a frame of reference based in the future.

The outcome is to create individuals who can programme their personal future—people who can create a goal so vivid and real that it becomes undeniable to the brain. This goal should be so powerful and compelling that it becomes easy to share the idea with others, and to convince them of its attraction.

Several questions can be considered by individuals when developing their own compelling future²:

- What specifically is wanted?
- Where am I now?
- What will the future look and feel like?
- How will I know when I have it?
- What does achieving the outcome give me?
- What resources do I have now, and what do I need in the future to achieve the outcome?
- What do I lose by achieving this future state, and what happens if I don't?

By simply going through this type of questioning process, the individual can develop a clear picture of what the next stage of the time line looks like. Furthermore, there is a tactile sense of what it feels, sounds and smells like. The desire is to create a sense of association between the here and now, and the future state, thus moving the time line from a theoretical picture to something that is real and internalised.

So, why should an organisation seek to release this potent force of power? Ghoshal and Bartlett³ suggest that, within every corporate hierarchy, numerous 'entrepreneurial hostages' are striving to break free—people who have tirelessly worked for the benefit of the company, doing the things that they were supposed to do, but never quite putting their own personal stamp on the activities. Once people realise the extent of the power, capability and influence they can attain, then it is possible to use this to help turn a good company into a leading-edge one.

Trait 2—Change Mastery

Consider for a moment all the recent management books, job adverts, and


training programmes that are being offered in the market. The one commonality is the requirement to find or train people able to manage change. This is now seen as a core competency for most firms, and is likely to be an ever increasing need, as the pace of change intensifies in the market place. However, the proposition is that the ability to 'manage' change will in itself be superseded. The requirement will be to master change, to shift from taking a reactive stance to a more proactive one. The difference might be analogous to the difference between a proficient pianist and one regarded as a virtuoso. The underlying skills and competencies might be broadly similar, but the difference is often around the performance, communication and the aura or sense of authority that surrounds the individual.

To become a master of change, people must first know the nature of change. There must be an acceptance that control, conformity and consistency are things of the past, and turbulence, trauma and transition are in effect the new normalised state. The first stage is to understand the common pattern that change follows, that of the Sigmoid curve.

The Sigmoid curve

Charles Handy suggests that by the time you know where to go, it's too late to get there⁴. People spend so much time in life working towards a point in the future, but failing to realise the benefits of the effort. Handy paints the picture of the Sigmoid curve (Figure 2), where there are natural up and down sides to all change processes. People typically climb the hill on the ascending side of

Figure 3—Stacked Sigmoid curve






Figure 2—Single Sigmoid curve

the curve, only to lose control and feel themselves fall over the other side.

The trick is to avoid the gravitational pull of the decline, by always keeping on the upside of the slope (Figure 3). People need to master the change, and build in a renewal process, whereby they kick off into an ascendancy phase just before moving into the decay stage of the cycle.

However, the problem is often in identifying the point where the new curve needs to be kick started. At what stage should a life, career or job change take place? Often this has to happen when things look fine, the income is OK and the knowledge bank is well stocked. This is because the future can be created only by breaking with the present. Sometimes it is important to shatter the existing paradigm that has been so successful, and create a new one for the future—a view that ignores the existing boundaries and limitations, and constructs a new set of rules by which to operate. This is often easier than most people think, because many of the current limitations are self-imposed, or readily accepted from another person or the organisation. People just have to take the courage to step out of the comfort zone, and go through the same process they undertook at the start of the earlier curve, and the one before that.

The secret for individuals who wish to have a sense of purpose and control over their lives is to see their future as a succession of curves, or hills to climb. Imagine waves hitting the beach, they go on and on relentlessly, each one looking the same, but in reality each one is individual and builds on the previous wave.

To operate like this, individuals must have a mindset that is able to work on parallel principles—on one hand to accept the current position in life and reap the just rewards, but on the other to assume that the current position is close to the peak. As with the sleeping cat that always appears to be partially awake, the successful individuals must always have one eye open, ready to jump start a change in their lives.

The art of re-energising people's careers is defined by building criteria by which individuals can map their position against the Sigmoid curve. Doing this in a market place that is both complex and confusing can be daunting and challenging. However, the only alternative to such an approach is to avoid taking any action, and thus reinforcing the idea of stagnation and eventual wastage. The following questions might be considered on a regular basis:

- What competencies do I have, and how do they align with the needs of the organisation?
- How do my competencies align with the needs of the market?
- What changes do I see coming in the market place over the next one, three and five years, and how marketable will I be if the changes come into play?
- How many years have I been in this current role, and how does this compare to colleagues and friends?
- What is the average cycle time for job changes in the industry, and how does this compare to my time in the role?
- If I need to make a transitional change, have I the necessary resources and energy to make the shift?
- Finally, do I have the willpower and courage to make such a transformation?

Individuals who regularly ask these questions will energise the organisation, as well as themselves. This constant use of inquiry, invention, and introspection creates a style of operating that enables new opportunities for the organisation. This might range from having advanced foresight about the market, through to developing a more innovative style of management. This benefit is even more pronounced when such style is assumed by a dominant leader within the organisation. Where the successful leader adopts this approach, so the organisation will generally follow, creating a bow wave of opportunities and successful outcomes in the market place.


Responses to change

In developing the capability to create and change one's own future, there is a supposition that a range of emotions will be stirred.

One model that neatly encapsulates many of the emotional change processes and feelings, is the U-loop (Figure 4). This is considered by Carnall⁵, where he suggests that change creates uncertainty, anxiety and stress. People will especially feel these emotions where the change process impacts upon their self-esteem, or position in the organisation. He suggests that there are five major stages in the change:

- **Denial** As a broad-brush suggestion, most people do not like change. The typical response to any shift in the environment is to think that it is OK to ignore it. People either do not believe it will

Figure 4—Change U-model



affect them, or are resigned to the idea that there is nothing they can do, so it is best to ignore the threat.

- **Defence** This is where people might take action to try and prevent the change from taking place. It might be simply withdrawing from the process, or, in some extremes, taking premeditated guerrilla action to stop the change from happening.
- **Discard** All change is about shifting a paradigm or view of the world. To move forward, one has to discard past behaviours, feelings and consequences—to let go, and be happy with what has gone.
- **Adapt** At this point, the past has been left behind, so people can begin to explore the new ways of working. This can be a confusing time, as the old and new are still likely to be in conflict with each other. However, with the correct support processes in place, the individual can quickly shift to the next stage of commitment.
- **Internalise** Once the change has been accepted, then the individual can adopt it as the norm, and internalise it as the current way of operating. At this point the energy returns, the process is owned, and the individual will, in many cases, try to convince others of its benefits.

As the rate of organisational change increases, there is a greater possibility that key people will be at the position of discard or adapt. In this position, the individual becomes less effective, so will that part of the organisation, and importantly the process of learning is likely to be impaired.

An effective organisation will ensure that it has processes in place that will respond to and counter any potential damage for the learning process. The options include:

- surfacing the U-model—make it okay for people to feel anxious and unhappy about change;
- creating space in the organisation where people can talk about their feelings and concerns;
- helping people to go through the U-process at a speed that is right for them, and effective for the organisation; and
- surfacing the hidden and unspoken issues that help fuel the grapevine. In making things public they will seem less of a problem.

If an organisation accepts that people need to go through such a transformation, then helping it to flow in an open and visible way can only help the learning process in the long term.

Becoming a change agent

Once an individual has attained a sense of change mastery, the shift must be to use this productively by becoming a change agent for the business. The definition of an agent for change is an individual that can 'effect' new and innovative ways of working across the organisation. The characterisation of this type of approach will depend upon a multitude of factors. However, the key things to consider are as follows:

- **Break-with view** Being able to take a new perspective on issues, and to do so by re-framing any personal viewpoints. In many cases an individual may have an established set of personal styles, behaviours, and operational maps that have served well over a long time. The need to break with these personal views will often be painful and traumatic. However, in flexing the 'break with' muscle, it will get stronger over time, reaching a point where, as an almost daily exercise, the individual will question and reflect on existing thoughts and maps of the world.

However, this break-with behaviour must be clearly understood in terms of the context, and situation where it needs to occur. In most cases, the 'break with' style is not simply a case of re-framing one's own view of the world. To stimulate learning across the organisation, individuals must be prepared to help others break with their view of the world. Such a step means that the change agent will have to overcome the resisting forces of tradition, social pressure and political reaction. This resisting pressure is often regarded as the cultural immune system. These are the unclear, and often hidden, forces that combine to prevent a change of practice from being accepted and implemented.

- **Loving turbulence** To facilitate and sell change effectively, the individual must also be able to adapt to the internal forces that drive the organisation. It is crucial to recognise that these forces do not conform to the more traditional idea of predictability and stability. The nature of organisations is now recognised as a model that is based upon unpredictability, complex relationships, organic structures, and boundaryless composition. The key attribute for any individual who seeks to share knowledge and wisdom throughout the organisation is to be gloriously happy when working with uncertainty. Every step that change agents take must be with the absolute knowledge that they operate in a world that is strange and bizarre. As such, when taking action designed to produce a specific outcome, unexpected responses are to be anticipated, and welcomed. They must be embraced and not seen as counter-productive to the change process.
- **Selling the change—the use of positics** Change agents will fail to stimulate learning and adaptation if they only operate within

The development of people who have mastered the art of change will ultimately help to create a climate of innovation, internal excitement, and creativity.

the confines of their domains. To be effective, the person needs to be able to take new ideas and sell them throughout the organisation. The individual must approach the internal sell with as much cunning⁶ and vigour as might be invested in selling a product to an external customer.

The change agent must employ all the familiar tools and techniques that are used by the sales professionals. The first step (1) is to map the territory, to determine the gatekeepers and supporters who will have an impact on the new idea. To do this, the individual must be able to target and understand the existing political relationships within the organisation, and, through this, to establish and maintain the appropriate contacts within the organisation, and so have immediate and current knowledge of the key stakeholders. (2) Being able to frame the idea in a way that is favourable to the various stakeholders is essential to the idea of selling across the business. Since an apple to one person will be a pear to another, the change agent must be able to select the words and phrases that mean something to the people who need to be influenced. (3) Key to building this influencing role is the ability to take a helicopter perspective, and stand back from the immediate projects, tasks and issues. Once the awareness and interest in the idea is created, the final step (4) is to manage the idea into the formal processes and systems of the organisation. The key to achieving this is to use the skill of negotiating. This is bargaining, where the outcome is a win for both parties, as opposed to a confrontational style, where the aim is just to force the idea or innovation into the organisation. The goal is to deliver an outcome that is reached efficiently and amicably, one that is built upon the idea of shared objectives, trust and synergistic solutions.

The essential core element that underpins each of these four issues, is the notion of *backstaging*⁷. Along with




Figure 5—Front-stage, backstage model

the more overt and public actions associated with managing change (project management, communication, training etc), there is a need for 'backstage activity' (Figure 5). This is typically the use of power skills to make interventions in the political and social systems. This is because learning and change are not always formal processes that can be managed simply and efficiently. They are often untidy, interrelated and overlapping processes that operate with shifting goals and team membership. As such, the processes that need to be used to underpin any adoption must also be loose in nature; hence the need to use backstage techniques that operate behind the formal levels of the organisation. The danger is that this might be viewed as a machiavellian approach. This is not so; it just requires a positive re-framing of the understanding of politics—to re-frame the action consciously, and think of *positic* actions, the 'positive application of politics'.

Ultimately, the goal behind 'selling the idea' is to mobilise the desired change across the organisation. Simply originating a new idea, or learning something that is of value to the business, is of little benefit if it is only used by a single person. The goal behind organisational learning, is the idea of enabling, creating, and sustaining learning beyond the range of the individual. Sometimes this will not happen on its own; hence a degree of influencing and push might be required to make it happen.

Seeing change as security

In achieving this state of change mastery, a new form of comfort and protection surfaces. Rosabeth Moss Kanter⁸ suggests that rather than

being a threatening and disrupting process, change can be regarded as a new form of security.

In the more traditional organisation, security is based upon acquisition of power—the power to hold resources, discipline people, control the finances and build personal empires. The suggestion is that in the learning organisation, people take a revised view of the world, where ability to master change rather than the accumulation of power gives long-term security.

The development of people who have mastered the art of change will ultimately help to create a climate of innovation, internal excitement, and creativity. As an end goal, creating a company full of people who have mastered change creates an organisation totally at ease with all forms of turbulence and distortion in the market place. In essence, it helps to create a truly adaptive and responsive organisation.

Trait 3—Continual Learning

If one looks back over the past 100 years, people's time lines have been built up of several stages. First there are the early years, where schooling takes precedence and the goal is to learn the basic life and academic skills that get the individual either into college or work. After this, further professional or vocational training might take place to enable the individual to perform an effective role in the business. In the final years, local training programmes might be used to 'top-up' the skills of the individual.

This model of the world is no longer appropriate, as life-long employment or professionalism is effectively obsolete. The rapid changes in society, and across industry, have revolutionised most markets and industries. Economic globalisation, changes in technology, demographic upheavals, and the high level of structural unemployment have placed people in a new and demanding position. No longer is it simply a case

of building a set of skills and competencies that will enable people to find a secure job. They must think along the lines of portfolios of opportunity, to ensure that personal revenue streams will be available 5, 10 and 30 years hence.

The most significant personal investment that a person can make is to create and maintain the capability to learn. So often, people view the early part of their years as the learning phase, and the middle and later years as the time when the rewards are reaped from the early investment. Now, it is no longer appropriate to view the process of learning as a one-off event. It needs to be a life-long way of operating. Each individual must 'learn to learn', and use this as a way to maintain their desired quality of life. To be competitive, individuals must have a rate of learning equal to, or faster than, that of the market in which they choose to operate. Take for




Figure 6—Learning opportunities

they must be seen to integrate as entirely complementary pursuits. The conceptual and practical separation of learning from the daily work activities is crazy, creates false barriers, and is wasteful. Honey and Mumford offer the view that learning can be viewed as being either conscious or unconscious, and on the job or off the job (Figure 6).

The proposition is that all four quadrants have value to the organi-

The most significant personal investment that a person can make is to create and maintain the capability to learn.

example the rate of change in the microprocessor market, or the software industry. Here, an individual must constantly re-educate in the new techniques and processes, simply to stay in the game.

The idea of personal learning has moved from a position where it was viewed as something an individual might do at evening class to something that is key for the organisation. An organisation that does not invest in the learning capability of its people is effectively bringing upon itself a process of death by a thousand cuts. The organisation will not see the effect of non-learning in the short term, it will simply fail to pick up on market trends, and slowly cease to create new ideas and products for future markets.

Instead of considering work and learning as alternative activities,

sation in its desire to create a learning organisation. The crucial step is when the individual can turn from looking at life and work in terms of tasks and activities, and can see each area as a learning opportunity or experience. When people start to reflect on and consider the activities they undertake, there can be a sudden realisation that things they have perceived as normal experiences have been wonderful opportunities to learn. When people talk about learning, they typically take the mental leap to the conscious/work quadrant, the natural response being to attend a new development programme, take part in a re-engineering workshop or undergo some new psychometric test. This model offers the view that there are three other areas of learning that are woefully forgotten and under-used:

- **Unconscious/Work** The opportunities that people have to learn at work are endless; for example, the changes in working practices that result from a reorganisation; the chance to become involved on a quality team; helping a new recruit; or taking time out to work with a colleague on a project team. All of these help to generate new skills, attitudes and behaviours and help the individual to develop new personal competencies.

- **Conscious/Social** In many cases, people split their work and social lives, maintaining that they are entirely different, and that the learning from one area is not transferable to another. So where people have undertaken development outside the work place, there should be encouragement to use this learning within the organisation. Consider the football manager who has attended a workshop on team management and motivation, the woman who is attending art class, and the man who is taking a course in oriental cooking. These are of immense use to the organisation when considered in a synergistic way; for example, the football manager using the coaching techniques with his team, the artist using creative techniques in a product development group, or the man considering products that the company might offer to niche areas in the retail food trade.

- **Unconscious/Social** Lastly, in many cases it is the unconscious social interaction that helps people to learn the most. If one considers any organisation, the basic glue that holds everything together is the interrelationship between the individuals. The organisation is essentially a social body that is held in place by the need to manufacture a product or service, and thus create an income. On this basis, anything that an individual learns about managing relationships outside the workplace should automatically be

adopted and used within the organisation. If one considers motivating a family member to undertake the chores, managing inter-family conflict, or running a social club, all of these social skills are of immense benefit to help running an effective organisation.

Conscious competence

The self-perception of change and learning can be nebulous and difficult to measure. Often, someone will take the decision to learn a new subject, only to slip back into the old style of behaviour once it gets difficult. The reason for this slippage is often grounded in the failure to understand the process of learning, and how it affects them personally.

Consider learning to play the guitar. If someone had never tried to play a musical instrument, then they would not be aware of the fact that they could not play (*unconscious incompetence*, Figure 7). The first time that person picks up a guitar, tries a few chords, and comes away with sore fingers, there is a clear realisation of how bad he or she is (*conscious incompetence*). After persevering and taking lessons, the musician emerges—someone who can play chords and notes with some ease and without too much pain (*conscious competence*). Finally, the time comes when the first few chords are mastered, and it is possible to play endless renditions of American Pie for friends and family. At this point, the process is happening, without the person having to pay detailed

attention to where every finger is placed (*unconscious competence*).

For the individual who wants to play solo (in a band or in the organisation), the need to identify, understand and welcome the stages of this cycle are paramount. It must be viewed as a process that is painful, but to be welcomed as part of the self-development journey. As the individual exercises the change and learning muscle, so it becomes easier every time and the process of learning to learn becomes the established. It is this inner awareness of the change and learning cycle that helps to create the sense of personal authority so important in developing people who will drive and underpin organisational-wide learning.

Generative learning

There are typically two types of learning within an organisation—*adaptive* and *generative*⁹. Adaptive learning, also known as *single-loop*¹⁰ (Figure 8), focuses on the delivery of a short-term response (knee jerk) to an emergent problem. The generative path, known as *double-loop* learning, focuses on re-aligning the problem. The generative style is effectively seen as 'learning how to learn', with the notion that this will be the key determinant for the development of learning within the organisation.

With adaptive, or single loop, learning, any change in the process or system is perceived at a superficial level. The underlying norms, values and assumptions are not questioned, or reflected upon in any way. For

example, if sales of a product are falling, the response might be to alter the promotional discount, increase the number of outlets or, in extreme cases, sack the product manager. Alternatively, where a team member is failing to meet the expected goals, the response might be to increase the level of monitoring, reduce the person's pay rating or, in many cases, take steps to discipline the person. In neither of these cases has any of the deep underlying issues been considered.

In taking a generative approach with these examples, deeper questions might be asked regarding the falling sales of a product. To what extent is there conflict in the market place with other products that the company sells? Are cross-portfolio problems emerging, has research and development lagged, so that the product is becoming dated, or are the sales performance targets at odds with the needs of the product turnover levels? For the individual who is under-performing, is there a problem at home causing concern, does the individual need a change in roles to make the work more interesting, or is there any conflict in the relationship with the manager?

The key to double-loop learning is in the art of turning the question back on the questioner. What underlying assumptions are being made? What are the values of the questioner in suggesting that a problem exists in the first place? For example, instead of the manager asking why the performance of a team member has fallen, he or she might ask why it took so long to notice, or what steps have been taken to develop the individuals in the team.

In helping an individual to move from an adaptive to a generative style of learning, the need is to migrate towards a style of self-reflection, inquiry and dialogue. In questioning the values and assumptions that underpin a situation, the first step is for the individual to have a clear and objective appreciation of his or her drivers and barriers, to appreciate why he or she might feel

Figure 7—Four-stage learning ladder





Figure 8—Single-, double- and triple-loop learning model



The challenge ethos is necessary because learning or change generally fails to occur unless the existing frame of reference is tested.

that something is a problem in the first place, and be able to consider the situation from an objective perspective, to question, examine and delve deeper into issues, and to find out what the underlying drivers are that have created the situation. However, it is important that this is undertaken from a positive stance, and not from a negative bias, where the goal is to find fault or blame others.

Generative learning often surfaces when a range of different options has been tried, but failed to bring about any long-term significant change; for example, where a product promotion or re-design failed to increase sales, or sending a team member on a number of courses failed to result in a change in personal performance. The enlightenment that can occur as a result of this exercise can be both rewarding and daunting. The rewarding parts are clear in that any change that causes deep levels of inquiry must be good for the organisation. The daunting part starts when one considers some of the discomfort that must be encountered to bring about the change. Of all the

most difficult things to do, moving down the level of insight requires that the question 'Why' is asked repeatedly (Panel 1).

In addition to creating a deep long-term solution to the problem, undergoing such a deep level of inquiry opens up a clearer insight in to the working of the organisation. Where people are prepared to accept the opportunities to develop collective self-knowledge, such a process can help to pull many of the old skeletons out of the cupboard. This leads to the third, but lesser known, form of inquiry, called *triple-loop learning*.

In triple-loop learning, the embedded customs, ethos and behaviours are held up for scrutiny. The organisation will consider the basic bedrock that it has been built upon, and its appropriateness. This approach forces the organisation to ask: 'What type of company do we want to be?', 'Do we wish to exist in this market?', and 'Who are the key stakeholders for the business?'.

Trait 4—Challenge Ethos

The challenge ethos is necessary because learning or change generally fails to occur unless the existing frame of reference is tested. One way in which this can happen is for individuals consciously to challenge the existing status quo, to make what might appear to be unreasonable or impractical demands on themselves, others and the incumbent systems.

Experimentation

The ability of an organisation to survive and prosper in the future is not defined and measured by the current level of knowledge. It is a function of individuals' ability to 'learn how to learn'—to question how people have assimilated what they know already, what processes were used, and how this experience can be understood, to increase the rate of learning within the organisation; fundamentally, taking experiences that have occurred, and turning them

into concrete learning concepts and generalisations for the future.

The proposition is that organisations and individuals cannot just sit back and wait for these experiences to occur. People must raise their consciousness, and create opportunities for learning.

Two key actions should take place. First, people should budget time to create learning opportunities, and learn from their experiences. Time should be set aside at the end of meetings to reflect upon, and learn from things that might have happened within the group. Second, the learning process is often improved by drawing on the ideas of variety and diversity. It is important to ensure that a range of experiences and outlooks is considered, so that just one frame of reference does not end up dominating the process. When this happens, the long-term effectiveness of the learning experience can be diminished as only one person's world view is used to drive the outcome from the group. From this, the proposition is that the most effective learning systems are those that tolerate, and actively welcome, diverse viewpoints.

Allow time and space to do things wrong

Over recent years, there has been a shift towards exerting greater levels of control and measurement within organisations. It is not untypical to find a manager who might have between 20–30 different measurement criteria that must be met and reported against, often on a daily basis. This need is driven by a number of factors, including the increasing use of cross-company benchmarking, and the abundance of quality control programmes. The resulting shift in mindset has often been to reduce the acceptable level of risk, and a bias against the acceptance of mistakes. So here lies the contradiction, the need to increase the levels of quality within a reducing cost range, while having to create increasing levels of innova-

Panel 1

The product sales are dropping. *Why?*

The sales force are not convincing the customer of the benefits of the product. *Why?*

The sales team did not all attend the product launch programme. *Why?*

They felt that it was not worth while wasting time on a course when they could be out with the customers. *Why?*

The feedback from the sales people who went on the early course was that it was of little benefit. *Why?*

The programme design was not piloted properly. *Why?*

Because of a last minute re-allocation of product development funds from training to the sales bonus plan.

Action: Review product development budgeting process.

Figure 9—Goshal and Bartlett contract

tion and invention within the organisation. The supposition is that the balance is possibly too far in the control direction, and that this has resulted in the failure to give people the time and space to do things wrong.


John Cleese, tells a wonderful story about 'Gordon the Guided Missile'. He suggests that, when a missile is first launched, it has a broad idea of where it is heading, but the absolute route is not known. What it does is continually to take a fix on where it is, what the current trajectory is and how far away it is from the intended goal, and then to modify the flight path accordingly. So, in essence, its progress builds on making many continuous mistakes. The key is that they are acceptable and known to contribute to the end goal, and minor short-term deviations are of little concern. It is the short-term minor deviations that tend to act as the spur for creativity and change within an organisation. Creative and innovative people stumble, trip and make horrible mistakes. Their highest and most acclaimed successes are often constructed on the rubble of humiliating failures.

In moving to this mindset, people must be able to fail creatively, and have the resilience to persist in the face of failure. The end result is that any creative mistakes will be treated as a learning opportunity¹¹.

Trait 5—Contract Alignment

It is important to appreciate that the six components in the soloist part of the model are all interrelated. The first four areas (compelling future, change mastery, challenge ethos and continual learning) are dependent on the fact that the individual has a sense of personal security. For many people, this sense of security comes in the contract with their organisation, the basic agreement that guarantees people the income and resources to satisfy their needs.

However, the recent changes in the make up of most industries have changed the nature of the contract. The



old picture of a stable market, with ever-increasing wage cycles, and guaranteed jobs, has effectively been ravaged. This change has manifested itself in two areas. The first and most obvious is the change to the legal side of the contract. The increase in part-time work, changing demographics, short-term contracts and the increase in self-employed people have made the legal process of contracting more important for the individual. However, it is the second and softer area that proves to be most important for the idea of organisational learning. This is the psychological contract, the unwritten and emotional bond that ties people and the organisation together.

During the past decade, the psychological contract has been modified. Company after company have pursued efficiencies through re-engineering, downsizing, quality programmes and outsourcing, and these have effectively ended the idea of life-long employment³. This has broken many of the shared beliefs that existed between the individual and the organisation. In essence, the nature of the long-term employee has undergone a radical restructuring, moving from predictability to turbulence and from long-term to short-term.

The question is how to turn this transformation to the advantage of the organisation and the individual. How can a new type of relationship be created that has the appropriate flexibility, but with the deeper psychological support that is essential to enable cross-organisation learning?

Rethinking the relationship

The suggestion is to move from a position where the organisation has the moral responsibility to provide life-long job security to one where it guarantees to provide life-long learning opportunities. In doing this, there is a reversal in role for the

senior managers from assuring that the company's competitiveness will be there to pay the wages to ensuring that the employees entrepreneurial initiative will safeguard their employability (Figure 9).

In considering this approach, Ghoshal and Bartlett³ raise a number of key points:

- The goal is for the employer to provide employees with continuous skill updating, to protect and enhance their personal flexibility, both inside and outside the company.
- The company must not abdicate responsibility to create a shared and embracing culture.
- People will move from being an asset that is exploited, to one where value is infused to enhance the competitiveness of the company.
- It requires people to move from a position of 'I am my job', to one where they take a personal responsibility for their learning capability.
- The shift towards empowerment means that people at the front line will make and take decisions, without waiting for top-level approval. Hence a process of internal 'letting go' will have to take place.
- The new contract is not a carte blanche for the individual to demand training and development at will. There must still be a form of criteria that assures that business benefit is ensured.
- The aim is to create a sense of 'shared success', a future where both parties knowingly add value

Where the individual and organisation can surface and realise their shared goals, there is likely to be greater sense of honesty, commitment and effort on the part of both parties.

to the relationship, so that joint benefits are accrued.

In helping individuals to create their sense of empowerment, the idea is to make the implicit explicit; that is, to take the issues that were once unspoken, and to move them to a position where they can be discussed in an open and honest relationship.

The new open contract

In moving the implicit issues into the explicit arena, the primary focus is in raising awareness of personal development as a factor within the contract. This element might be summarised as the *meta skills*¹² components. An example of the type of things that might be covered include:

- learning balance,
- portfolio of opportunities,
- explicit trust description,
- motivation drivers, and
- violation clause.

This is not meant to be a definitive list of things to include in the new open contract. It does, however, surface some of the more hidden and personal issues that are not normally considered within the contracting process. The ultimate idea is to bring into place a sense of alignment (Figure 10), helping the organisation, and its people, to realise that prosperity is built upon the notion of shared, rather than selfish, success. Where the individual and organisa-

tion can surface and realise their shared goals, there is likely to be greater sense of honesty, commitment and effort on the part of both parties.

- *Learning Balance* The idea is to make the notion of shared learning objectives an explicit part of the contract. An individual might attend a training programme focused primarily on personal development, but the business would still be aware of what commercial benefits might be accrued in the long-term. Alternatively, the individual might focus on developing skills that are explicitly for the business to take advantage of, but there would still be a process by which the individual can leverage some personal gain out of the experience. The proposition is not to take a selfish or short-term view, but to focus on the shared destiny, and seek to derive benefit from a synergistic relationship.

- *Portfolio of opportunities* One can (with some degree of confidence) assume that an individual will not do the same job, in the same place, with the same company for his or her whole career. Hence it is important that opportunities are taken to build a breadth of experience and knowledge, as well as the depth that most people attempt to attain. The contract should highlight what opportunities will be open to the individual, and how they can be used to develop the individual's personal career.

- *Explicit trust description* If one considers all the various activities, relationships and projects that occur within an organisation, they are all dependent upon the balance of trust between people. The level of trust is like the backbone in an organisation, it provides the underlying strength to ensure that promises are kept, and work is completed on time.

There might be a great deal of benefit in including a trust index within the contract; some form of indicator that explicitly highlights the balance of trust in the relationship. The difficulty comes when considering what the criteria might be for measuring the degree of trust that exist, but certain criteria might include:

—*Integrity.* The extent to which honesty and truthfulness are maintained.

—*Consistency.* The degree of reliability and predictability contained within the relationship.


—*Loyalty.* Benevolence, or the willingness to protect, support and encourage each other.

—*Openness.* Mental accessibility, or the willingness to share ideas and information freely¹³.

- *Motivation drivers* So often massive assumptions are made about what motivates an individual. The classic example might be where the organisation wants to praise and energise someone for something done at work. To do this, it is arranged for the individual to have lunch with a director or senior manager. However, the individual finds this extremely embarrassing and stressful, and the resulting action might be to demotivate the individual. Motivational drivers should be included as part of the contract and should be reviewed on a regular basis to consider their effectiveness.

- *Violation Clause* Possibly the final element needs to be a description of how any violation of the contract might be managed. Clearly, for the legal side of the contract, violation is simple (theft, redundancy, etc.), but for the softer elements, such clear-cut criteria might not be so simple. The

Figure 10—Shared success alignment



contract violation can take place in one of three ways:

- inadvertently, where both parties are able and willing, but an interruption takes place because of a misunderstanding;
- through disruption, where circumstances make it impossible for either one of the parties to fulfil their side of the contract; and
- by reneging, where one side of the contract makes a conscious effort to break the agreement.

Although it might be possible to define exact ways and means by which the psychological contract will be violated, it becomes impractical to specify all possible infringements. It might be more valuable to spend time discussing and agreeing ways in which potential violations can be rectified, before the idea of full separation occurs.

If one takes the time to look beyond all the issues raised about the psychological contract, the most important idea is to make the undiscussible discussible by taking those issues that are normally left to happenstance, and talking them through to a point where there is a more open relationship. This relationship is based upon helping organisations and people to better understand their needs, and the options for meeting these. It is taking the often unspoken idea of 'What's in it for me?', and turning it into 'What's in it for us?'.

Trait 6—Considerate Nature

As with any change process focused on enhancing the role of individuals, the danger is in making the pendulum swing too far. The goal is to give people the opportunity to gain a sense of empowerment, alignment and confidence, without pushing them into a position where they become selfish and

insensitive. To use any of the positive points raised in the first five areas of the soloist stage, people must be able to work effectively with other people. To do this they must understand how to map the balance of their relationship with others, how to understand when the relationship is working, or alternatively, when things are less positive.

An effective tool for doing this is offered by Dr. Stephen Covey in his book, *The Seven Habits of Highly Effective People*¹⁴. He suggests that the nature of a relationship between people can be symbolised through the use of an 'emotional bank account'. It is a way of indicating the negative and positive factors that interplay between people. An individual can reflect on a relationship with another person, and consider whether he or she is in credit or debit with that other person. By debit, this means

*people must be
able to work
effectively with
other people*

that the individual might have been discourteous, untruthful or negative towards another. By credit, it might be a case of helping, talking to or sharing information with the other person.

Although only a simple model, it has the effect of being a powerful metaphor for mapping relationships. It gives individuals a process whereby they can understand what consideration they might have paid to other people's feelings or needs. It also helps to reflect on action that might be taken to get the bank account back into balance again.

Conclusion

Learning starts with the individual. Without people who are committed, energised, competent, and happy in their workplace, then little new knowledge will be generated.

The trick then is to ensure that processes are in place to facilitate a culture where people are encouraged and rewarded to learn at an individual level and to focus on creating the personal traits of:


- compelling future,
- change mastery,
- continual learning,
- challenging ethos,
- contract alignment, and
- considerate nature.

It is important to note that the six traits are achievable. There is nothing new or innovative; they are simply a distillation and representation of many existing ideas and practices. The secret is in their unification, and the resulting synergy that surfaces once they are in place (Figure 11).

Although these six traits might be viewed as falling into the softer territory of management, the realisation must be that without sufficient focus on this area, the long-term viability of the organisation can be put at risk. Just consider the companies that have disappeared over the past 25 years, as they failed to keep in touch with the rapid transformations in the market place!

The key to maintaining a competitive presence in future markets will be through the generation, management,

Figure 11—Six-trait integration




and utilisation of knowledge. With the recent changes that have spread across the communications market, the economic and production power lies more with a company's intellectual and service capabilities than with its hard assets, such as land, plant or equipment. The value that is embedded within most products is derived from the knowledge-based intangibles that people add; for example, technological know-how, product design, marketing presentation etc. Hence, in an industry and society based upon learning and knowledge, the knowledge worker is the organisation's single greatest asset¹⁵.

This point is so clearly made by Peter Senge¹⁶ where he suggests that: '*Organisations learn only through individuals who learn. Individual learning does not guarantee organisational learning, but without it, no organisational learning occurs*'.

All organisations must therefore ask themselves: 'How many people do we have who are able to take a solo role at any time?', 'What strategy do we have for enhancing the capability to learn?', 'How can we effectively use our knowledge as capital?', and 'How can we enhance the alignment between the three areas?'

To return to the musical analogy, consider you have just become the proud owner of a priceless antique Stradivarius. Who would you trust to play it, Yehudi Menuhin, or a dispirited and ill-prepared student? Who then do you want to perform in your organisation?



References

- 1 COLLINS and PORRAS. Building your Company Vision. *Harvard Business Review*, Sept/Oct 1996, p. 66.
- 2 JAMES and WOODSMALL. Time Line Therapy, p. 79.
- 3 GHOSHAL and BARTLETT, The Individualised Corporation. Heinemann, 1998.
- 4 HANDY, CHARLES. The Empty Raincoat, p. 49.
- 5 CARNALL, COLIN. Managing Change for Organisations. Prentice Hall, 1995.
- 6 BUCHANAN and BODY. The Expertise of the Change Agent, p. 75. Prentice Hall, 1992.
- 7 FAIRHURST and SARR. The Art of Framing. Josey Bass, 1996.
- 8 KANTER, ROSABETH. The Change Masters. Unwin Paperbacks, 1983.
- 9 SCHEIN, ED. Organisational and Managerial Culture as a Facilitator or Inhibitor of Organisational Learning. 1994. <http://learning.mit.edu/res/wp/10004.html>
- 10 ARGYRIS and SCHON. Organisational Learning II. Addison OD Series, 1996.
- 11 GIOIA. Contrasts and Convergences in Creativity. In FORD and GIOIA (Ed.), Creative Action in Organisations, p. 317. 1995.
- 12 DE MEUSE and TORNOW. <http://deming.eng.clemson.edu>
- 13 HOSMER. Trust the Connecting Link Between Organisational Theory and Philosophical Ethics. *Academy of Management Review*, 1995, 20, 379.

14 COVEY, STEPHEN. Seven Habits of Highly Effective People. Simon and Schuster, 1992.

15 NONAKA and TAKEUCHI, The Knowledge Creating Company. Oxford University Press, 1995.

16 SENGE, PETER. The Fifth Discipline. Century Business Books, 1990.

Biography



Mick Cope
BT Strategy and
Business Services

Mick Cope completed his MBA in 1994, and is currently undertaking a Doctorate in Business Administration in Organisational Learning. He joined BT in 1974, and rapidly established himself as a particularly uninspiring apprentice, with the words still ringing in his ears 'you'll never become a Technical Officer, lad'. After spending a number of years on the CSS project, he moved into the organisation development field, where he supported the introduction of Total Quality into Northern London District. After this he managed the project office for ISO9000 implementation in Worldwide Networks. He subsequently managed the BT pricing process and worked as a global pricing analyst. He is currently organisational development manager for the Strategy and Business Services division. Contact number 01268 786036 and fax 01268 783705.

Closure of Cardiff Empire Exchange

Since 1912, telephony service to Cardiff has been provided from a site adjacent to Cardiff Arms Park, the national rugby stadium. When the Welsh Rugby Union decided to construct a new stadium, in which to host the 1999 rugby World Cup, BT agreed to vacate its Cardiff Empire exchange, which would be demolished to make way for the new Millennium Stadium. This article tells the story of the engineering work involved in clearing Empire exchange.

Introduction

Cardiff Empire exchange (Figure 1), located in Park Street, dates from the introduction of automatic telephony in Cardiff shortly after the Second World War. Building work, which had started in 1938 but was halted by the war, was completed in 1947, and the exchange was opened by the Postmaster General of the day, Earl de la Ware, on 16 February 1952. At its height, the exchange provided service to over 40 000 town centre customers using Strowger technology, and contained a manual exchange of over 300 controlling and enquiry positions.

To cope with the growth in trunk traffic in the 1960s, a major building extension was completed in 1967. At that stage the exchange housed a group switching centre (GSC), and later a transit network main switching centre (MSC). The continued growth of trunk traffic in the 1970s necessitated the opening of a second

GSC unit, Cardiff Stadium, in the adjacent Stadium House, and so the Strowger unit was designated *Cardiff Empire* (after the nearby Empire pool, built for the 1958 Empire Games). In 1988, as part of the switch modernisation programme, customers were transferred to a System X unit in Stadium House, and the GSC function ceased in 1990, with the closure of the analogue trunk network.

When, on 23 February 1996, the Millennium Commission† gave its assent to a proposal by the Welsh Rugby Union (WRU) for funding of £46M to build an all-year all-weather stadium in time for the 1999 rugby World Cup, the clock started ticking for BT to clear Cardiff Empire exchange by the end of 1997.

† The Millennium Commission, established under the National Lottery Act (1993), receives 20% of the proceeds of the Lottery until 31 December 2000.

Figure 1 – Cardiff Empire exchange (Stadium House is on the right)




Figure 2—Directory Assistance interconnection

Realising the extent of the work involved, assessed at over 100 000 man-hours, BT Networks and Systems set up a project management team to organise the clearance. An audit of the building revealed:

- a measurement and analysis centre;
- a Directory Assistance centre;
- a TV network switching centre;
- a frequency division multiplex (FDM) transmission node;
- an audio transmission area serving over 3000 private circuits;
- an extensive 'double-decker' main distribution frame (MDF);
- a redundant trunk test suite and an emergency switchboard;
- a cable chamber, which was to be retained, but extending in height to over a metre above ground level;
- an earth electrode system installed under the South Stand of the present rugby stadium; and
- a redundant diesel-operated engine set, complete with 35 000 litres of diesel!

Directory Assistance

Directory Assistance (DA), formerly Directory Enquiries, has come a long way from a handful of telephonists answering calls on key-and-lamp units, and looking up numbers in phonebooks or microfiche. DA voice traffic—dialling code '192'—is routed via the trunk network to the digital derived service network (DDSN), a network of fully interconnected 5ESS switches, also used for services such as 0800, 0345, LinkLine etc. Calls are routed (based on the originating exchange) to specific DDSN switches serving DA sectors; Cardiff, being in the Leicester sector,



is served from the Birmingham DDSN switch.

From there calls are routed over 2 Mbit/s links to a remote interface signalling unit (RISLU) in the DA centre, which distributes traffic to individual operators. Number information is held on the Directory Assistance system (DAS) database, access being provided via KiloStream circuits from a central DAS computer site to local terminal group controller (TGC) equipment, and thence to individual visual display units (VDUs). Once the called number is found, the call is transferred to an automatic voice response (AVR) unit, a fresh connection made from the AVR port to the DDSN switch, and the required number is transmitted from the AVR unit to the customer.

To relocate the DA centre meant preparing a new site from scratch. The ninth floor of Stadium House, formerly the original Cardiff Stadium GSC, was cleared of redundant crossbar GSC equipment, and then totally refurbished by BT Building Services to modern office standards. To avoid congestion on the existing lift facilities, a complete window section was removed on the ninth floor, and the building contractor installed an external hoist. For several months the floor resembled a building site in nature.

With 80 equipped positions, Cardiff is the largest DA centre in the UK. BT Operator Services had made it plain that there must be no break in service in effecting the changeover (other than the normal

23.00–07.30 hours overnight closure) and to cover this, new TGC, AVR and RISLU equipment was installed by contract in readiness. Owing to DDSN capacity restrictions, the additional RISLU could not be connected while the existing RISLU was in service, and so a comprehensive databuild and in-service test programme was run in the overnight period in the immediate two-week period prior to changeover.

To meet the requirements of the current display screen environment legislation, the computer terminals in the new DA centre are Windows NT-based personal computers (PCs), connected via a local area network (LAN) instead of directly connected VDU-type terminals. The LAN wiring terminates on shelf LAN cards in the TGCs, proving the versatility of the original DAS system and TGC design from the 1980s, now interworking with Windows NT PCs. The LAN is also interconnected to a Sun server, which holds the applications software and individual log-in details, and which will facilitate software upgrades. This is illustrated in Figure 2.

Testing went smoothly, revealing only a cross-connection wiring fault between the RISLU and the operator positions interface cabinet, and the changeover was effected unnoticed by the customers, as telephonist staff logged off at 23.00 hours at the old centre in the Empire House exchange, and logged on again at 07.30 hours the following morning in the new centre in Stadium House.

TV Network Switching Centre

The function of the TV network switching centre (TVNSC) is to switch, or distribute, TV vision and sound signals between studios and transmitter sites, as required by programme schedules, and to act as a specialised transmission area, where radio, music etc. feeds can be centralised.

'Specialised' proved to be an understatement, as on further investigation it transpired that the Cardiff TVNSC also included:

- the main distribution feed from UK Channel 4 to S4C (Sianel 4 Cymru),
- TVRO (receive-only) distribution racks feeding local customers via polyethylene quad copper (polyquad) cable, and
- remote control and supervisory racks for 6 GHz radio systems at the nearby Wenallt radio station.

In the Empire TVNSC the vision circuits were provided on a mixture of coaxial cable and optical fibre routings, with obsolescent terminal equipment, much of which dated from the 1970s. In planning the new TVNSC in conjunction with BT Broadcast Services, a decision was made to upgrade all the circuits onto digital line systems using new vision codecs.

Staff were relocated by BT Property to free up an apparatus area on the eighth floor of the adjacent Stadium House, which was then fitted out with ventilation plant and overhead ironwork, and the appropriate transmission hardware installed.

The hub of the new TVNSC is a suite of proprietary Video Product Group (VPG) vision codecs running at 140 Mbit/s, feeding transmitter sites and studios via optical fibre. Although internal fibre ties to the old TVNSC were available, it was

decided to use a non-standard changeover, to minimise customer downtime, and in effect allow the customer (for example, the BBC studios) to decide the timing of their changeover. This was done by opening the optical fibre joint nearest the exchange, diverting the spare fibres to the new TVNSC, and then routing the replacement feed over these fibres (Figure 3). Once the customer had accepted the new circuits, the fibre tail to the old TVNSC was recovered and the joint closed.

To provide audio services, dedicated cables ran to several local sites, with cable pressurisation fed from the old TVNSC. These were rationalised onto three cables, rejoined, pressurised from the Stadium House MDF area, and presented as copper tails into the new TVNSC. To minimise equalisation problems (for example, for music circuits), the existing 62-type amplifiers were 'lifted and shifted' to the new TVNSC. VHF stereo circuits were routed over Musicline 2000 equipment, which provides for up to six 15 kHz stereo circuits in either direction, plus capacity for data channels (depending on configuration).

Interestingly, the change to VPG technology enticed one of the TV companies to change over from 'sound in sync', where the audio component is transmitted by time division multiplex techniques within the vision synchronisation pulse, to

straight stereo channels, derived from the VPG codecs, at a commercial benefit to BT!

In recognition that the 'TV' in TVNSC stands for 'television', a local television presenter was enlisted to open formally the TVNSC on 4 June 1997, in order to publicise BT's role in the preparations for the new stadium.

Earth Electrode System

The adjacent Stadium House, built in 1978, was initially equipped with a total energy installation¹. Earthing for this was provided by a series of 4.5 m rods (reaching the old river bed level), located under the South Stand of the present rugby ground. In addition, to this was bonded the Cardiff cable tunnel² ironwork, and the original Empire exchange earth plate. This had produced a very-low-resistance earth-electrode system, with an as-installed measurement of 0.07 Ω . Although such a low reading could not be reproduced, the very low resistance was confirmed by one of the original BT designers, now retired, who lives locally.

It was clear that the earth mat under the South Stand would be destroyed in the construction of the new stadium. Although the historic needs for earthing had changed—the withdrawal of shared service and DC2 signalling systems, and the supersession of the total energy plant by a direct 11 kV feed—it was felt that,

Figure 3—TVNSC optical fibre diversion




Figure 4—Soil resistivity testing

with the importance of Stadium House as a processor node site, as low a figure as possible should be maintained.


A decision was taken to install the replacement earth mat in a strip of land being transferred from WRU to BT on the northern edge of Stadium House. Soil resistivity readings using the four-electrode method (Figure 4) indicated that over 150 conventional copper electrodes would be needed to achieve the target figure.

Other approaches were considered, therefore, and after extensive consultation with John Laing Construction Ltd, the WRU's contractor, and Laing's subcontractor, Omega Furse, the lighting conductor specialists, the solution was adopted whereby Laing would provide a replacement electrode using Marconite. Marconite is a black granular material, which is mixed with ordinary Portland cement. When used with copper electrodes, the Marconite not only adheres to the electrode, but also 'leeches' into the surrounding soil, reducing the overall resistivity.

Ten blocks of Marconite each of one cubic metre were sunk in the strip, and when connected gave a reading of 0.4Ω which, together with the other electrodes, brought the combined figure to 0.21Ω . This was regarded as acceptable, and the old earth mat under the South Stand was disconnected.

Analogue Circuit Offload

The audit had revealed over 3000 analogue audio circuits spread in roughly equal proportions over the inter-exchange junction copper cable network and the FDM network. Techniques differed for offloading, or upgrading, these. Generally the circuits on the FDM network were reprovided as 'analogue over digital' circuits over the KiloStream network. This work was carried out by BT Customer Network Services staff, as part of the national FDM offload program. The shorter-distance 'copper offload' circuits were



rejumpered at their exchange onto new routings, typically as individual 64 kbit/s channels in 2 Mbit/s multiplex equipment.

60 kHz pilot tone

The clearing of individual FDM groups and supergroups brought to light a potential problem in that Cardiff was a distribution station in the 60 kHz pilot tone distribution ring. The 60 kHz tone is used to synchronise oscillator frequencies in the channel translating equipment³, and it was necessary to reroute the feed to West Wales to maintain continuity when the Cardiff FDM site closed.

Branching panels

Branching panels⁴, containing passive audio components, are used when a private circuit has to be split into several 'branches'; a typical application is a data circuit to a building society computer, feeding several offices in the same locality, but spread over more than one serving exchange. Typically the long-distance link, if analogue, is

routed over the FDM network, and the local links over the junction network. This is illustrated in Figure 5.


While BT staff are adept at ceasing and providing private circuits, reprovding whole branch networks, while at the same time offloading them, was an untried scenario, and considerable thought was given as how best to effect such a changeover, without tying up significant resource at each point.

The solution adopted (Figure 5) was, once the long-distance link had been reprovided, to provide the old branching panel as 'spur 0' off the new panel. Individual spurs could then be upgraded separately off the new panel, with 'spur 0' ceased on completion.

DC paths

Some customer private circuits require a continuous DC metallic path for correct operation. Although a DC path, as a facility, was withdrawn for new circuits in 1990, over 50 circuits were found which required DC paths for signalling purposes. Discussions

Figure 5—Reparenting of branching panels



Cable Chamber

One of the most challenging aspects of the clearance was the cable chamber. Although Stadium House is served by a deep-level cable tunnel², local access network cables to the street cabinets in the immediate vicinity, and to the former Roath satellite exchange area, are routed via the Empire cable chamber. After considering the options for re-routing these, it was decided that the most appropriate course of action would be to retain the cable chamber, after the exchange proper was demolished. Any cables made redundant by the analogue offload work could simply be left *in situ*, or recovered at a later date.

The position changed considerably when, in July 1996, the WRU appointed its main construction contractor (John Laing Construction Ltd), and detailed design work on the new stadium commenced. From this it became clear that:

- the cable chamber would have to be reduced in height,
- the cable chamber would have to be reduced in width for most of its northern length,
- cables would have to be moved onto bearers on the south wall,
- the new soffit slab would have to take vehicular loadings, and
- the exchange manhole would have to be rebuilt.

In order to safeguard customer service, it was a BT requirement that the reconstruction of the cable chamber had to take place *before* the building was formally vacated and handed over; it was evident that this work was on the critical path for the clearance.

Plans and detailed method statements from Laing for the reconstruction were scrutinised carefully by BT and its consultants,

Percy Thomas Partnership (Architects) Ltd. BT requirements on the security aspects, gas detection, mechanical and electrical services, and drainage were drawn up. Although BT already had experience of retaining operational cable chambers beneath commercial developments, particularly in the City of London, this was the first known case where the 'integrity' of the cable chamber was affected, and so an additional requirement was the need for substantial protection to the cables while reconstruction was underway.


After the plans for the reconstructed chamber had been approved, work was set in hand on the phased recovery of the cables being made redundant through the clearance. Progress was interrupted for six weeks in late 1996 while asbestos-based fire-stopping material, which had been disturbed by contract staff, was removed. Recovery of the final cables was carefully coordinated with the closing stages of the analogue offload, and the chamber was handed

over to Laing as soon as the offload was completed.

Cable bearers and ironwork from the nearby BT Enterprise House, a former exchange building, were renovated and reinstalled by Laing's staff on the chamber south walls, with temporary bearers around one corner pilaster, which could not be removed for structural reasons. BT staff moved the remaining cables from the north wall onto the new bearers, and Laings then put up a substantial wooden protective hoarding, which had also been designed to allow emergency access by BT jointers, if necessary (see Figure 7). After the remaining ironwork had been removed, shuttering formwork was erected, and concrete poured, using the chamber's original air vents; it seemed as if they had been designed with this future use in mind!

Although the Empire cable chamber interconnects with Stadium House, provision was also made for a new entrance into the chamber from the exchange yard, to ensure maximum

Figure 7—Reconstruction of cable chamber (by permission of John Laing Construction Ltd.)



flexibility for the future, and to ensure an additional escape route, in case of emergency. In addition, a footway cover opening was built into the soffit slab to facilitate future cabling operations, and anchor irons, of the appropriate strength for winching operations, were cast into the floor and walls.

Because the area in the immediate vicinity of Empire House is to be rebuilt in a piazza-style development, and this will necessitate partial reconstruction of the exchange west-end manhole, a duct seal has been built into the west wall of the new chamber in readiness. The rebuilding of the manhole, and the removal, and making good, of the corner pilaster, will take place after the demolition of the exchange building in the summer of 1998.

Acknowledgement

The author would like to acknowledge the support of the many Networks and Systems engineering staff who worked with him on this project; in particular, Bill Winkley, and later Paul Shackley, for their assistance on the project management work, Clive Webb of Percy Thomas Partnership (Architects) Ltd, who advised on the civil engineering aspects, and Laing's design engineers, who handled the detailed requirements for the cable chamber.

References

- 1 JOHNSON, N. G. The Total Energy Installation at the Cardiff Telephone Exchange. *Post Off. Electr. Eng. J.*, July 1980, **73**, p. 92.
- 2 JAMES, N. A. Cardiff Cable Tunnel. *Post Off. Electr. Eng. J.*, Oct. 1980, **73**, p. 177.
- 3 BLACKHALL, R. W. E. FDM Equipment. *Post Off. Electr. Eng. J.*, April 1973, **66**, p. 38.
- 4 LAWRENCE, A. J. Audio Transmission Equipment. *Post Off. Electr. Eng. J.*, April 1977, **70**, p. 29.

Biography



Michael Clemitson
BT Networks and
Systems

Michael Clemitson joined the then GPO in 1968 as a direct management entrant. Based in Preston Telephone Area, he was employed on exchange design, and later on telecommunications consultancy work. Moving to Cardiff on promotion in 1980, he had operational responsibility for the Wales and Marches Regional Prestel Centre, later moving to the planning division, where, on the District reorganisation, he was responsible for the District's switch modernisation programme. Prior to taking charge of the Cardiff Empire clearance project, he was responsible for the Wales and West wideband provision, circuit provision control and ISDN 30 provision functions.

This quarter, Peter Cochrane turns his attention to virtual reality and looks at how it has moved from the laboratory into the real world. No longer is it solely the prerogative of scientists, playing games in the laboratory, rather it is now a vital part of the business world.

A Toy No Longer

For many engineers, designers, scientists and medics virtual reality has already become a mainstream technology. Much of it is a far cry from the popular hype and vision of head-mounted displays and gloves first seen in the games environment. It is far more pragmatic! Screen-based virtual reality now dominates by virtue of the fact that it is readily available, easy-to-use and realises great benefit. Today, engineering design is probably the dominant application area for it with medicine uses coming a very close second.

Designing all but the simplest artefacts on paper has had its day. Moving to the N -dimensional world of virtual reality provides greater clarity and understanding, but more importantly great savings are to be had. The need to shorten drastically the time to market, and get products right first time, have brought this technology to the fore. Producing anything from a mobile telephone through to an automobile or an airline terminal can see savings of over 30 per cent in time and money through the use of virtual reality for the design. Similar savings can also be realised in process design where logistics and components are critical factors in production flow and control.

Visualising the final design in full (virtual) operation is a vital step in getting it right, but without incurring the expense of an actual build. All the best organisations are deep into this area as a matter of course. It is no longer a laboratory toy, it is an essential tool.

Medical applications are developing rapidly with everything from body fly-throughs to operation simulations and animations. But in this area, along with many others, it

is the mixing of the real and virtual worlds where the greatest advantage probably lies. Combining telepresence and virtual reality allows surgeons to benefit from a real-world view augmented by computer generated simulations. In recent trials, surgeons have been able to stand 'one inside the other' at a distance to experience new surgical techniques for the first time, or receive reassurance during a first-time solo operation. This technology is equally applicable to the repair of a personal computer, jet engine or heating plant. It offers a new and alternative approach to education—a metaphoric guide on the inside!

Managing a modern company can be like flying a 747 airliner with 100 fold the instrumentation than is actually necessary. It is not uncommon to be data rich and information poor. You can see the temperature of the toilet seats, but your altitude and heading is anybody's guess. Here virtual reality has a very big, and largely unrecognised, part to play. It is ideally suited to the representation of highly-complex and data-rich situations. Beneath this very thin veneer of civilisation we are natural



of science, mathematics and many other topics. It is principally a medium for direct experience and soon we will have the ability to step inside the atom or the molecule, fly a proton and experience fission, rather than just gaze at a set of complex equations. For the first time we will see and feel the binding energies in the alignment process of a long-chain molecule while simultaneously viewing the equations and associated graphical information. For many well understood systems and situations we can already view and handle mathematical functions and models in a new way. They no longer have to be artificially frozen in time and

Experiments have seen board members digest 2 Gbytes of information in a matter of 20 seconds.....


hunters. One of our primary skills is visual correlation. We did not evolve the skills for gazing at spread sheets. But put data into an N -dimensional virtual reality field and the results are stunning. Experiments have seen board members digest 2 Gbytes of information in a matter of 20 seconds—and understand a complex situation for the first time. This is more than the entire contents of the Encyclopaedia Britannica. The trick is animated graphics, not spreadsheets and static two-dimensional graphs, as well as moving, interactive three-dimensional colour.

Virtual reality also offers significant potential for the teaching

space by the limitation of the paper page, but can be alive with N -dimensional interactivity.

It is interesting to reflect that only 50 years ago classes at schools and universities were commonly augmented by practical demonstrations on a laboratory bench that may still be in the front row today. Effectively that was virtual reality 50 years ago—you just sat and watched someone else do it. Today much more can be done on the screen by everyone. It may be the tool we have been looking for—instant education and understanding, just in time! The question is—will anyone be in the real classroom or will it all go virtual?

Software Projects— Evolutionary versus Big- Bang Delivery by Felix Redmill



The author opens his book with the statement that information technology (IT) projects are inherently risky. Through a swift walk through the history of software development, ample justification is made for this opening statement. Good project management is, of course, all about managing this risk. Success means that software is delivered on time, within cost, and to the user's real requirements. How this can be achieved forms the essential question that this book addresses.

The book moves on to standard models used for software project management. It discusses the waterfall, the vee, the matrix and the spiral models for software development. The discussion is rigorous and the author obviously has a deep understanding of his subject. However, in chapter four the author declares that a good project infrastructure is the most important aspect of project management but does not mention any of the standard methods, such as Prince or Telstar. Later, he develops some generic ideas about the requirements for a project infrastructure which, although well-founded, are insufficiently developed to put straight into practice.

However the author is really addressing the fundamental questions of the best software-development model to adopt to succeed in software-

project management. In chapter six, the concept of evolutionary development (ED) is presented and explained. He defines this as the provision of functionality of a system in a number of deliveries over a period of time. The required system becomes something that grows to eventually become a complete solution. According to the author it will grow in three ways. Firstly, functions are added as in the original specification. Secondly, changes are made as a result of feedback and communication from the users. Finally, new requirements are delivered which come about because the user, through experience gained on the initial releases, is at last able to articulate his/her real needs. Changes can be to any of the delivered functions, including the first. The key idea is that ED will deliver eventually what is really required, through an evolution to full functionality.

The author is far too experienced to present this as the panacea to all problems and admits that an ED project can still be late and over budget. It will, however, be a system that meets the real needs and wishes of the user. This concept could be dismissed as being the same as simple de-scoping and phasing of a project but the author convincingly shows that this is not the case. Through reasoned argument, he says that the approach is new and can yield significant advantages. What is less convincing is the way that he dismisses rapid application development (RAD) as being different because it deals only with short-timescale user-based systems rather than longer strategically-based systems. This is a provocative statement that would cause some debate among the RAD practitioners.

The rest of the book explores project management and how it differs when using ED. The author does not attempt to hide the fact that the approach brings its own problems but faces up to them and demonstrates that they can be overcome. The importance of company culture is not ignored and the author illustrates from personal experience how the required changes can be achieved.


Overall, the book is extremely well written and easy to read. It is systematically structured with chapters alternating between theoretical concepts and case histories and each chapter includes a summary of the main points. This is an excellent book for seasoned project managers who are looking for new ideas or want to stand back and think about their own approach to project management. It is also a good book for those just starting out, as the case histories and the theoretical content amply demonstrate the challenges of the job. It also clearly sets out the breadth of skills required to be a successful project manager.

*Published by John Wiley & Sons Ltd.
£34.95. xvii + 254 pp.
ISBN 0-471-933430*

Reviewed by David Bason

Signalling in ATM Networks

by Raif O. Onvural and Rao
Cherukuri



Asynchronous transfer mode (ATM) technology is being positioned and developed to support a wide range of emerging multimedia applications and to integrate all types of networking services and applications. It was chosen by the International Telecommunication Union (ITU) as the underlying transport technology for the broadband integrated services digital network (B-ISDN) and as such ATM has been the subject of extensive standardisation activity, not only within ITU but also within the

ATM Forum (ATM-F). A good deal of this standardisation work has concentrated upon the signalling systems and protocols which provide users of ATM networks, and the B-ISDN, with the ability to establish, modify and release, on demand, virtual channel and virtual path connections. The initial signalling standards produced by both ATM-F and ITU were based upon those already devised for circuit switched 64 kbit/s narrowband networks with additions and modifications to take account of the traffic, connectivity and quality-of-service (QoS) aspects associated with ATM networks.

This book describes the basic initial signalling standards, principally produced by the ATM-F, for use at the user-network interface (UNI) and network-node interface (NNI) for the setup and release of point-to-point and point-to-multipoint ATM connections. These ATM-F UNI and NNI standards described are known as *UNI Version 3.1/4.0* and the *private network-network node interface (P-NNI)* respectively. The equivalent ITU UNI standards are also covered, along with the ATM-F variants, known as Q.2931 and Q.2971. The UNI protocol descriptions are supplemented with example signalling system behaviours, messages types and message contents needed to support services such as LAN emulation over ATM, video-on-demand (VoD) and IP over ATM. In addition, the operation of the signalling ATM adaptation layer (SAAL) is described which provides a reliable underlying signalling message transport function.

Although the ATM-T and ITU UNI protocols are described together, there are differences between the two 'equivalents' which are not always made clear. These differences, arising mainly from the way ATM traffic and QoS has been defined within ATM-F and ITU, may leave the reader thinking that ATM-F and ITU-T UNI signalling protocols are interchangeable. In practise, this is not the case without appropriate interworking functions which are not described in the book. Also notable by its absence is a description of the ITU's C7-based

NNI signalling protocol, the broadband integrated services user part (B-ISUP). This NNI protocol is now available on larger ATM switches and is likely to be of interest to public network operators (PNOs) such as BT.

In addition to the above basic broadband signalling standards, there is also a description of the signalling standards defined to support B-ISDN supplementary services; that is, calling line identity presentation (CLIP), direct dialling in (DDI) etc. Following this is a chapter on interworking ATM networks with narrowband and frame-relay networks, from both a signalling and transport perspective. The ATM layer transfer capabilities (ATCs) defined by both the ATM-F and the ITU receive a detailed treatment in a separate chapter together with how they are supported by the signalling protocols. The explanation of the ATCs is supplemented with details of how the associated traffic parameters can be negotiated, using additional signalling procedures and messages, during and after connection establishment.

Overall, it is a very readable book, giving a valuable insight into ATM and B-ISDN signalling standards that are now approved and on the table today. It contains explanations and background information not available within the raw standards which would be especially beneficial to someone new to this area. The book is, however, focused too much towards private ATM networks, with emphasis on the signalling standards produced by the ATM Forum. I would have liked therefore to have seen a more balanced exposé, describing additional signalling systems, also of relevance to PNOs, such as ITU's B-ISUP and European Telecommunications Standards Institute (ETSI) VB5.2 protocol for controlling broadband, or combined broadband and narrowband, access networks.

Finally, there is only a brief mention in the book of what future B-ISDN signalling capabilities are in the pipeline. These capabilities move broadband signalling systems away from their narrowband origin towards a target solution using separated call- and bearer-signalling

mechanisms. This change of approach is needed to support future fixed mobile convergence concepts and advanced B-ISDN multiparty, multimedia Conferencing services.

Published by Artech House Publishers.


£56.00. xviii + 337 pp.

ISBN 0-89006-871-2

Reviewed by Roger Morley

Quality of Service in Telecommunications

by A. P. Oodan, K. E. Ward and A. W. Mullee



If you had to reach out for a book on quality of service in telecommunications you would have been hard pressed to find one until this multi-authored work was published. That unique selling position should guarantee it a place in reference libraries and on the shelves of those involved in regulation and standards setting in the telecommunications industry.

While the work is tightly focused on telecommunications it does cover a wide field and seeks to incorporate elements of marketing, economics and operations management to legitimately introduce the customer, shareholder and operator perspectives.

The multi-author approach leads to stylistic variations and some uneven treatment of subjects but is an ambitious work and the end-of-chapter references are very good for those readers wishing to pursue an issue in greater depth.

Proliferation of new service providers, network providers and the ever widening range of services in support of an information society is creating confusion in what was never coherently measured in quality of service terms outside of the PTTs and even then with little reference to true customer requirements. Therefore, the logical approach by the authors to what is quality, where are we now, what are the issues and what roles have to be discharged is helpful. The architectural framework for measuring overall quality and disaggregating components is also logical and represents a sound basis for a regulator or an international standards body to approach new services.

From an operational perspective, the middle section is particularly good in that it covers the service surround and the related service management, network-performance engineering, network integration and the economics of quality comprehensively. Without drilling too deep in any single subject it exposes the key issues and trade-offs.

In an increasingly diverse world, product performance (from a customer perspective) will be critical to network providers who have to become expert at managing customer value and in the words of Bradley T. Gale, 'that means creating quality of service that customers can see'. Keeping pace with what customers value will be a tricky act in the fast moving world of telecommunications, and regulators may well be forced to restrict their interest in quality-of-service measurements to the commodity level.

Overall the authors have achieved their object in filling a gap in the published work in this area and their proposed framework is an ideal basic structure. However, it is in the interests of the service and network providers to distort their products' marketing mix to service particular market sectors and it would be a brave regulator who strayed into the softer areas of product specification and customer satisfaction. Equally, it may become unrealistic for network providers to hold to commodity product specifications in a market

situation where customers are using bought-in applications to move up the value chain and find either the basic network-product specification or its reliability to be inadequate for the new business applications.

These issues are teased out in the various models which suggest that the authors' proposed framework for quality measurement has a reasonable shelf life.


Published by The Institution of Electrical Engineers.

£55.00. xxiii + 375 pp.

ISBN 0-85296-919-8

Reviewed by Tom Denniss

Introduction to Radio Propagation for Fixed and Mobile Communications *by John Doble*



This book is part of the Artech House Mobile Communications series which covers all aspects of mobile communications. The author aims to give the reader an insight into the propagation effects encountered when working on point-to-point fixed links (using microwave line-of-sight transmission) and the problems encountered when designing a successful mobile network.

The book is divided into two sections, the first deals exclusively with fixed links and the second section is devoted to mobile communications and the major differences between fixed and mobile links.

The first section takes the reader into all the important aspects of propagation and covers the back-

ground to microwave line-of-sight links. The role of the atmosphere and anomalous propagation conditions are covered which helps explain how the refractive index is affected by both the height above sea level and atmospheric conditions. The effects of aeriels are covered and there is discussion as to the best compromise between antenna gain and beam-width. All the topics covered are related to typical frequencies and hop length for microwave links in the United Kingdom. The planner is also alerted to the fact that a route has to work in all weather conditions and calculations for worst month conditions are included.

There are references made to simplified atmospheric models and are not just the International Telecommunications Union—Radiocommunications Section (ITU-R) models but also models for regional variations within the United Kingdom. This helps provide more detailed references for transmission-network planners. The disadvantages of microwave propagation are also covered including detailed explanations of how signals can be affected by fading. The mechanics of fading are covered and there are worked examples of how signal fading can be calculated and discussion of how this will affect analogue and digital signals. There is a complete chapter devoted to diversity reception and how this can be used to overcome the effects of fading. The worked examples in the book, particularly in the chapter on system planning, help to ensure that full understanding of all the principles, covered in the section on fixed links, is achieved.

The second section of the book has to cover a wide range of services and a large section of the radio spectrum as mobile communications can take place anywhere from very high frequency (30–300 MHz) to extremely high frequency (30–300 GHz). The different aspects covered are biased towards speech links but mention is made of such services as paging and all aspects of mobile communication from private mobile radio to analogue and digital cellular services.

The propagation of very high frequency (VHF) and ultra high

frequency (UHF) signals are covered with more than adequate detail. In addition, there are notes on the effects of the profile of the ground and buildings within built-up areas. I was pleased to note that an entire chapter is devoted to communications into and within buildings. A point is made of emphasising the effects of building materials such as tinted glass and foil-backed plaster board. There is also a mention of how to predict signal loss in an office building at 900 MHz using experimental results. It would have

been nice to have seen a similar propagation model at 1.8 GHz for global system for mobile (GSM) services although references are made to previously published results. The reuse of frequencies is also covered and with the growing demands on radio spectrum due to the growth of communications services this will be a very important consideration for planners in the future.

The book gives a great deal of information in a very down-to-earth manner on how to plan transmission

links for both fixed- and mobile-communications systems. The author is to be congratulated for the way in which a very complex field of communications has been presented. I feel that this book would be a very important addition to any radio system planner's library.

Published by Artech House Publishers.

£55.00. xviii + 189 pp.

ISBN 0-89006-529-2

Reviewed by Robert Snary

telecom focus

Two More IRIDIUM Satellites in Orbit

On 8 December 1997, two more IRIDIUM satellites were successfully launched, bringing the total number of IRIDIUM satellites in orbit to 41. These are part of IRIDIUM's 66-satellite wireless personal telecommunications network designed to offer full global coverage and a variety of communications services including voice, data, fax and paging.

'With this launch in China, we've set a new record by deploying 41 satellites in eight months from three countries,' said Christopher Galvin, Motorola's Chief Executive Officer.

BT Makes Multi-Million Pound Investment for Internet and Data Network of the Future

BT has announced a £300 million network modernisation programme to support the rapid growth in Internet and data communications services. These services are expected to account for about 90 per cent of all corporate communications traffic by the year 2003. This investment will be used to develop and expand BT's existing high-speed, high-capacity broadband multi-service network. BT already has the biggest network of this kind in Europe, offering high-speed services to the finance, retail, IT, education, manufacturing and

media sectors. The announcement further demonstrates BT's commitment to be the world's leading provider of data and broadband services into the next millennium. Chris Earnshaw, Managing Director of BT's Networks and Systems, said, 'Our goal is simple—to help our customers to take advantage of the data wave and the tremendous business opportunities offered by the Internet and company intranets.

'With this investment, not only can we meet demand—we can exceed it. It will enable BT to continue leading the implementation of a national broadband infrastructure for the integration of voice, data and video applications over a single network.'

BT will create a common multi-service broadband network to meet demand by business and residential customers up to the millennium and beyond. It will use the latest asynchronous transfer mode (ATM) technology, and is an expansion of BT's high-speed CellStream service—the UK's first national ATM service. ATM is a key technology for future integrated multi-service communications as it is designed to handle, in a flexible, cost-effective and scaleable manner, large volumes of Internet (IP) traffic and multimedia applications which are a mix of voice, data and image.

BT has gained invaluable ATM experience since launching the CellStream service in June 1996, and now plans to invest in growing and

developing that network to support its existing and future broadband communications services. Among the first services to be migrated onto this multi-service network will be BT's first broadband data service—switched multi-megabit data service (SMDS)—to be followed by Frame Relay as well as BT's range of Internet products and services. By delivering many products and services from a single managed-network infrastructure, BT will be able to respond more quickly to changes in customer demand between the existing products as well as reducing its 'time to market' for future services and capabilities.

BT's investment clearly signals its intention to lead the world into the twenty-first century by offering state-of-the-art data services as well as traditional voice services. BT expects to see enormous demand for integrated broadband services which is underlined by the fact that the UK data market alone is expected to be worth more than £7.5 billion by the year 2003, from £3.5 billion today. Chris Earnshaw added, 'BT's investment programme is testament to its commitment to provide its customers with the applications and network services to flourish in the new information age.

'The management and exploitation of information and data is key to competitive advantage. With high-speed services and applications from BT, the commercial opportunities for our customers are unlimited.'

BT Trial Set to Transform Home Telephone Line into Digital Highway

BT has announced that it is to trial Home Highway—a new digital communications service that will transform a customer's existing home telephone line into a high-speed Internet and multimedia connection. The trial, starting early next year, will involve about 300 households in the Midlands.

BT Home Highway will use the latest digital-access technology and operate over existing BT telephone lines at speeds several times faster than the fastest modem. The technology offers the customer a choice of using a combination of analogue or digital connections. These are: two digital data channels, one analogue voice channel and one digital data channel or just one analogue voice channel.

This means a family, for example, could surf the Internet at high speeds and make and receive telephone calls simultaneously. Currently, users with a single analogue line have to disconnect from the Internet to use the telephone. It also means that customers can benefit from a fast and clear digital connection without having to give up their existing telephone number, equipment and range of BT Select Services.

Other uses of BT Home Highway include:

- a much faster way of accessing the Internet than using a conventional analogue telephone line and modem;
- the ability to view goods on the Internet using a digital-data channel order over the spare telephone channel;
- an effective means of transmitting data from home while still keeping a channel free for personal use;
- high-speed access to educational services such as BT Home-Campus—developed to provide on-line learning; and

- improved experience of multi-player Internet-based computer games by allowing users to talk on the telephone while playing.

Rupert Gavin, managing director of BT's Consumer Division, said, 'The aim of the BT Home Highway trial is to assess customers' experience and enjoyment of using the Internet and on-line entertainment services by providing a faster, more affordable and reliable connection to the home.'

'The service will open the doors to a number of new applications as well as enhancing existing telephone services. We are still working on the functionality and packaging and will be testing a number of pricing options during the trial to find out what customers want and how much they are willing to pay.'

All customers will be recruited by targeted marketing activity to allow BT to achieve an appropriate demographic sample of the UK population.

BT Home Highway is designed to be easy-to-install, providing an additional socket with two digital ports and two analogue ports connected to a customer's existing telephone equipment. Customers simply plug-in their existing telephone equipment and computers in much the same way as they do today.

BT Welcomes OfTel Fair Trading Decision

BT welcomed the decision by the Director General of Telecommunications not to take formal enforcement action on BT's Chargecard business under the Fair Trading Condition in BT's licence. The Director General had issued a Provisional Order in July 1997 to stop BT subsidising its Chargecards from other areas of its business. Before issuing a Final Order he sought the advice of the newly-formed Advisory Body on Fair Trading in Telecommunications which has concluded that BT was not in breach of the Fair Trading Condition. A spokesman for BT said, 'We are pleased that the Advisory Body concluded that BT has not abused its market position. This is the first time

this procedure has been used, and we are pleased at the regulator's decision to seek the Advisory Body's advice and to follow it.'

BT and ESB to Offer Irish Communications

BT and the Republic of Ireland company Electricity Supply Board (ESB) have announced that they have agreed in principle to form a joint venture to offer communications services in Ireland, western Europe's fastest growing economy. The new company will be a 50-50 joint venture between BT and ESBI Telecommunications, ESB's telecommunications investment unit. Investment by the alliance will be approximately £75 million in the first five years growing to £130 million over 10 years and it is expected that the company will have more than 300 personnel.

The joint venture will immediately have access to ESB's digital-microwave network and will use the company's overhead electricity cables and ducts to expand the network. This network, one of the most extensive in the country, will be linked to the BT system.

Headquartered in Dublin, the company plans to offer a range of communications products and services to Ireland's business community. Building on the success of BT's current operation in the Republic of Ireland and Northern Ireland, the company will link business parks in large urban areas and provide international products from Concert Communications Services. Other services will be introduced when full liberalisation of the Irish communications markets occurs in January 2000.

Sir Peter Bonfield, BT's chief executive, said, 'We are delighted to have formed this alliance with ESB. The new company will benefit greatly from ESB's high reputation and knowledge of the market.'

'The Republic of Ireland is a major base for multinational companies and call centres which already know the advantages of using Concert services. We must be present in this market to fulfil our ambitions to be the number one pan-European operator.'

Ken O'Hara, ESB's chief executive, said, 'The Irish telecommunications market is growing rapidly and opportunities are emerging as the market is liberalised. The new ESB/BT joint venture provides ESB with a unique and innovative opportunity to invest in this market.'

'BT is a highly reputable, progressive group which can offer global coverage and is one of the best partners for ESB in this market.'

Lucy Woods, chief executive of BT Northern Ireland and the Republic of Ireland, added, 'While we could have grown organically, this alliance with ESB will give the joint venture a major opportunity to benefit from the track record of both organisations.'

OVUM Report Highlights Dangers of High Interconnection Charges

Competition in the telecommunications market is in danger of being obstructed by high and diverse interconnection charges as identified in a new report researched and published by Ovum. The report, commissioned by BT, examines interconnection charges—the rates paid by telecommunications companies to interconnect their networks and switching centres with those of other national operators. A BT spokesman said, 'Interconnection is a critical issue affecting competition as national telecommunications markets open up. Unless clear and fair interconnection rates are offered, new entrants in the telecommunications markets are subject to high charges by the incumbent operators. This results in an unfair burden on the new entrant and endangers the development of genuine competition.'

The report examines the interconnect prices for call-termination services across 12 countries and 17 operators and reveals a wide disparity in the prices operators charge.

BT to License Gilbert Scott Payphone Kiosks

BT has announced that it plans to license its trademark Gilbert Scott K6

payphone kiosk for use by competitors. The move is to promote competition where planning requirements by local authorities prevent other operators siting their modern kiosks. Licences will be issued on condition that the operators use a colour other than red, and that they make it clear that the service is provided by an operator other than BT.

Until now, BT has been reluctant to allow use of the K6 design by other operators because of its strong association with the company and because the kiosk provides restricted access for disabled or infirm customers. However, because of the continued insistence by a few local authorities that only the K6 can be used in designated areas, and the failure of an earlier site swapping agreement to open up these areas to competition, BT has decided to grant licences to help promote customer choice.

BT Operations Director Bob Warner said, 'We want our customers to have a choice of payphone provider—but then to choose to use a BT Payphone. Today's announcement that we are licensing the Gilbert Scott kiosk underlines our commitment to fair competition.'

AT&T and TCG to Merge

AT&T has announced that it has signed a definitive merger agreement with Teleport Communications Group (TCG) for an all-stock transaction valued at approximately \$11.3 billion.

AT&T said the merger, expected to close in mid-to-late 1998, will accelerate its efforts to bring end-to-end communications services to American businesses.

'This is a great match with powerful financial and strategic synergies for both companies,' said AT&T Chairman Michael Armstrong. 'Joining forces with TCG will speed AT&T's entry into the local business market, reduce our costs and enable us to provide businesses the any-distance services they want.'

'TCG has more fibre route miles and serves more businesses in more cities than any other competitive local-services company,' said Armstrong.

TCG will become the foundation on which AT&T will build a new unit with accountability for the company's local services. TCG also will continue to provide wholesale services.

Consumer Choice Must Come First

Oftel has stressed the need for consumer choice and wide competition in the new information age and has published the first part of its evidence to a Parliamentary inquiry into developments in audio-visual communications and the regulation of broadcasting.

Oftel's initial evidence to the Department of Culture Media and Select Committee concentrates on the technological background and policy issues which it believes Government should decide as it paves the way for UK consumers and industry to get the most out of revolutionary developments in broadcasting, telecommunications and information technology, including the Internet.

The second part of its evidence, to be presented to the Committee as soon as possible, will concentrate on how the converged communications market could be regulated.

Publishing the first part of Oftel's evidence, Don Cruickshank, Director General of Telecommunications said, 'Developments in the world of communications are moving so fast that the Select Committee's inquiry into the implications for both consumers and industry is both welcome and timely. The European Union has already launched a debate into the competition and regulation issues caused by the rapid breaking down of barriers between broadcasting, telecommunications and information technology. The outcome of that debate could be new communications laws which bite on the UK. We must decide quickly what outcomes we want, so the UK we can play a full part in shaping the outcome of that debate.'

'The Select Committee has set out to explore a number of questions, ranging from the nature of technological changes affecting audio-visual communications and broadcasting, to

the impact of those technological changes on the structure of communications regulation in the UK. Oftel will be responding to these questions—and raise some more—in two separate submissions.

'In the first part of our evidence, which we have now lodged with the Committee, we deal with the impact of technological change, and stress the importance of deciding what the policy objectives should be. Oftel's view is that communications technology is changing so fast, that the UK's policy for exploiting the opportunities must be flexible enough to deal with future change that as yet we can't even begin to imagine.

'We are on the brink of a world where consumers will receive moving pictures, information, entertainment and voice services in a variety of new and different ways. Oftel's prime aim is to see the development of maximum consumer choice with best value for money.

'These rapid developments raise questions about if, why and how the converged communications market should be regulated. Increasingly, developments in one regulator's area raise issues for other regulators. This has the effect of blurring boundaries which could lead to inconsistency in decisions and the risk of 'double jeopardy' for companies wishing to invest in information age services.

'Once Government is clear what its policy objectives are, and what special rules are needed, it can then decide the best way to enforce them.'

Oftel Sets Out Future Approach to Competition Issues

The Director General of Telecommunications, Don Cruickshank, has set out guidelines on how Oftel will go about responding to competition issues in the UK telecommunications market.

Don Cruickshank said, 'Competition analysis is an integral part of Oftel's work to achieve its goal of providing the best deal for customers in terms of quality, choice and value.

'However, I do not propose to introduce rigid, formal tests for

dealing with competition issues. I want to keep the flexibility to analyse competition issues on a case-by-case basis. The framework I have set out is very similar to the approach used by other competition authorities, such as the Monopolies and Mergers Commission, the Office of Fair Trading and the European Commission.'

Helping Coventry Area to Meet the Challenges of the New Millennium

Oftel has announced important new moves to help ensure the Coventry area can meet the hi-tech challenges of the new millennium.

As Coventry has expanded and new businesses have set up in the area, demand for telephone numbers has rocketed. That's why Coventry is to become the sixth area to have its code changed to the new 02 range.

From the year 2000 Coventry's code will move from 01203 to 024. The new code offers a tenfold increase in the supply of numbers in the Coventry area, from 1 million to 10 million. The 024 proposal is in response to expansion such as:

- an increased targeting of the Coventry area by telecommunications companies who are offering multiple numbers per home, through services such as distinctive ringing;
- increasing demand for multiple numbers in the home and workplace from advances such as Internet and PCs, faxes, and teleworking;
- the provision of direct-dial numbers to students in halls of residences at both of Coventry's expanding universities, and;
- continuing business growth.

Don Cruickshank said, 'Everyone is looking to our hi-tech future—telephone numbers are one of the keys to unlocking the gateway to the information superhighway.

'Failing to take action now could see the Coventry code area running

out of numbers very early in the next century.'

The area served by the present 01203 range includes Coventry, Nuneaton and Bedworth, Kenilworth, and parts of northern Warwickshire.

A typical number currently is: 01203-987654. After the year 2000 this number will change (by an addition of 024-76) to: 024-7698-7654.

The addition of 76 helps to ensure that there is room for further growth in the number range. This is a part of Oftel's ongoing review and monitoring of telephone number provision and demand.

National Grid for Learning

The Government should make sure that their plans for the National Grid For Learning allow all schools to have affordable access to high-speed networks, and a choice of a wide range of on-line educational packages, Don Cruickshank, Director General of Telecommunications said.

Publishing Oftel's response to the Government's consultation on its plans for a National Grid, Don Cruickshank said, 'I welcome this initiative as a great way of giving our pupils and teachers the chance to share fast access to high-quality information and services in the information age.

'I have recommended to the Government that they should adopt a framework which unbundles telecommunications links from other services (content, training, local-area networking and equipment within schools) needed for the Grid. They are already proposing that schools should be able to purchase telecommunications connections to Internet service providers separately. I am recommending that they go a stage further and unbundle the links that Internet service providers and publishers of content are using. Separate purchasing of telecommunications links will make it easier to ensure that schools have access to any content they wish to use.'

BRITISH TELECOMMUNICATIONS ENGINEERING

ISSN 0262-401X

Published in April, July, October and January by The Institution of British Telecommunications Engineers, Post Point G012, 8-10 Gresham Street, London, EC2V 7AG.

Formerly *The Post Office Electrical Engineers' Journal* Vols. 1-74: April 1908-January 1982.)

The Board of Editors is not responsible for any statements made nor the opinions expressed in any of the articles or correspondence in *British Telecommunications Engineering* or the *Structured Information Programme* unless any such statement is made specifically by the Board.

© 1998: The Institution of British Telecommunications Engineers.

British Telecommunications Engineering and the *Structured Information Programme* are printed in Great Britain by Headley Brothers Ltd.

Subscriptions and Back Numbers

Annual subscriptions (including postage)—companies, universities, libraries and other bodies £60.00 UK, £65.00 overseas; private individuals (paid by personal cheque) £36.00 UK, £41.00 overseas.

Single copies and back numbers—companies etc. £15.00 including postage UK, £16.25 including postage overseas; private individuals (paid by personal cheque) £9.00 including postage UK, £10.25 including postage overseas.

Overseas customers should pay by a sterling cheque drawn on a London bank, a sterling Eurocheque or a sterling travellers cheque payable in London.

Price to BT employees: call +44 171 356 8050 for details.

Orders, by post only, should be addressed to *British Telecommunications Engineering Journal* (Sales), PP: G012, 8-10 Gresham Street, London EC2V 7AG.

Remittances for all items should be made payable to 'IBTE' and should be crossed '& Co'.

Advertisements

All enquiries relating to advertisement space reservations should be addressed to The Advertisement Manager, *British Telecommunications Engineering Journal*, PP: G012, 8-10 Gresham Street, London EC2V 7AG. (Telephone: +44 171 356 8050. Fax: +44 171 356 7942.)

Communications

All communications should be addressed to the Editorial Office, *British Telecommunications Engineering Journal*, PP: G012, 8-10 Gresham Street, London EC2V 7AG. (Telephone: +44 171 356 8050. Fax: +44 171 356 7942.)

Copyright

The entire contents of *British Telecommunications Engineering* and the *Structured Information Programme* are covered by general copyright and the Board of Editors' permission is necessary for reprinting extracts.

Authorisation to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by The Institution of British Telecommunications Engineers for users registered with the Copyright Clearance Centre's (CCC's) Transactional Reporting Service, provided that the base fee of \$5.00 per copy is paid directly to CCC, 27 Congress Street, Salem, MA 01970, USA. For those organisations that have been granted a photocopy license by CCC, a separate system of payment has been arranged. Payment is additionally required for copying of articles published prior to 1978.

0262-401X/98 \$5.00 + .00



THE INSTITUTION OF BRITISH TELECOMMUNICATIONS ENGINEERS

IBTE Local Contacts

The following is a list of local contacts to whom enquiries about IBTE membership and services can be directed.

Zone / Centre	Contact	Telephone Number
London	Terry McCullough	(01707) 601151
Martlesham Heath	Peter Shaw (in Ipswich)	(01473) 224488
Midlands	Eric Pedersen	(0121) 230 4344
North East	Geoff Jenkinson	(01226) 248141
North West	David White	(0161) 600 2468
Northern Home Counties	Keith Setchell	(01234) 277937
Northern Ireland	Allan McDougall	(01232) 215301
Paris	IBTE Office	+44 171 356 8008
Scotland	Raymond MacMillan	0800 671991
South West	Peter Shaw (in Bristol)	(0117) 972 3216
Southern Home Counties	Nick Smith	(01227) 474823
Wales	Geoff Oatway	(01492) 545130

IBTE International HelpLine: +44 171 356 8008

IBTE On-Line (BT intranet): http://mhwww.boat.bt.co.uk/ibte_nat/ibtehome.htm

BRITISH TELECOMMUNICATIONS ENGINEERING

Journal Subscription Order Form

PLEASE COMPLETE THE FORM IN BLOCK CAPITALS AND RETURN IT, WITH THE APPROPRIATE SUBSCRIPTION PAYMENT, TO THE FOLLOWING ADDRESS:

The IBTE Administration Manager
Post Point G012
8-10 Gresham Street
London EC2V 7AG
United Kingdom

Annual subscription rates (four issues, including postage and packaging) are as follows:

Companies, universities, libraries and other bodies: **£60.00 UK; £65.00 Overseas**
Private individuals (to be paid for by personal cheque): **£36.00 UK; £41.00 Overseas**

Please supply four quarterly issues of *British Telecommunications Engineering*. I enclose a cheque/postal order to cover the yearly subscription as given above.

(Cheques and postal orders, payable to 'IBTE', should be enclosed with the order. Overseas customers should pay a sterling cheque drawn on a London bank, a sterling Eurocheque or a sterling travellers cheque payable in London. Alternatively, customers may wish to use a local subscription agent.)

Title (Mr/Mrs/Ms etc) _____ Initials _____ Surname _____

Mailing Address _____

Postcode (Zip Code) _____ Country _____

Signature _____

Date _____

Please state the issue you wish your subscription to commence: _____
(April, July, October or January)

VAT Registration Number (zero rate): 243 5917 54

Photocopies of this form are acceptable.



FEDERATION OF TELECOMMUNICATIONS ENGINEERS
OF THE EUROPEAN COMMUNITY (FITCE)



37th European Telecommunications Congress

Institute of Education, London, 24-28 August 1998

'Diverging Roles in a Converging Marketplace'

Telecommunications is converging with computing and broadcasting, opening new opportunities both for traditional operators and new arrivals. By the time of the Congress, the European telecommunications market will have been open for eight months to new entrants from all industries. The Congress will include papers exploring the technical and commercial opportunities that this situation offers.

The programme is yet to be finalised, but topics are likely to include the impact of new technology; new roles and challenges for telcos; differentiation in, and the structure of, the telecommunications business; life with regulations and the regulator; working with multiple standards and parallel solutions; increasing user expectations; price versus service; access, interconnection, service and network management in a multi-provider environment; and electronic commerce.

Further information, including details about how to register for the Congress, is available from national FITCE associations, the FITCE'98 Web site <http://www.btwebworld.com/fitce98/> or the Congress organising office, BT Conference Unit on +44 181 477 2842.

Sixth IEE Conference on 'Telecommunications'

**Edinburgh Conference Centre, Heriot-Watt University, Edinburgh, UK
29 March-1 April 1998**

The IEE Conference on Telecommunications is the premier professional meeting in the UK at which to review, discuss, debate and learn about telecommunications developments. Those who attend should come away with a sense of where telecommunications is heading worldwide. The content of the conference will interest researchers, engineers, managers, operators and regulators—in fact all who play a part in shaping, operating and developing telecommunications.

The objective of the conference is to be informative and to appeal to a wide range of interests: a conference where people can learn from experts in a particular field and where matters of the moment can be discussed freely among people knowledgeable in the subject matter.

The conference seeks to reflect and explore new developments, evolutionary trends and current best practice in telecommunications in terms of technology, networks, services and applications, and to relate these to associated business and regulatory issues. It will deal with all aspects of research, development, deployment and applications of telecommunications systems, equipment, networks and services and will embrace the broader operational, organisational, regional, regulatory and business issues inherent in telecommunications today.

The conference will be held at Edinburgh Conference Centre, Heriot-Watt University, Edinburgh, UK. The University is situated just six miles from the city centre in a delightful rural setting, offering modern conference facilities and comfortable accommodation.

Note that, as the IBTE is co-sponsoring this Conference, IBTE Members may register at the reduced Members' rate.

For further details, please contact the Secretariat: Janet McCready/Anna O'Kane, Conference Services, Institution of Electrical Engineers, Savoy Place, London WC2R 0BL; Tel: +44 171 344 5469/8425; Fax: +44 171 240 8830; e-mail: jmccready@iee.org.uk/aokane@iee.org.uk; Internet: <http://www.iee.org.uk/Conf/ICT/>

Contents

Guest Editorial: New Challenges, New Opportunities	265
Switzerland—Life in a New BT Joint Venture	266
BT WebWorld™—Web Hosting for Every Business Need	273
Games People Play—The Wireplay Experience	279
A High-Performance Network for Gaming—Wireplay	284
Local Area Mobility	290
The Underground and Copper Access Network Infrastructure	296
Fighting the Information Terrorists	300
Beyond the Handset: Audio for the Next Millennium	306
A Wearable Communication Space	311
Interacting Demands for Multimedia Services	318
Creating the Corporate Soloist Who plays the Stradivarius?	324
Closure of Cardiff Telephone Exchange	337

